10-2023

Does the consolidated feed matter?

Shihao YU

Singapore Management University, shihaoyu@smu.edu.sg

Does the Consolidated Feed Matter?[1]

Shihao Yu

October 31, 2023

(Job Market Paper)

[1]Shihao Yu, Columbia University, Department of Industrial Engineering and Operations Research, Email: sy3119@columbia.edu. I thank Albert J. Menkveld, Bart Yueshen Zhou, Björn Hagströmer (discussant), Sida Li, Carole Comerton-Forde, and Lucas Saru for their helpful suggestions and discussions. In addition, I thank participants at the Microstructure Exchange and the 2022 Plato MI3 Academic Conference for their helpful comments. I gratefully acknowledge the Netherlands Organization for Scientific Research (NWO) for a Research Talent grant. I acknowledge computing resources from Columbia University's Shared Research Computing Facility project.

(Job Market Paper)

Does the Consolidated Feed Matter?

Abstract

We examine the role of the consolidated feed in a fast and fragmented market. Utilizing an exogenous technological shock that significantly reduces the feed's latency in the US equities market, we find that surprisingly, a faster consolidated feed worsens market liquidity. We show it might stem from informed slow traders benefiting from the speed increase as evidenced by larger price impacts and higher execution efficiency. As a response, high-frequency market makers widen their spreads to compensate for greater adverse selection costs. Additionally, when the consolidated feed is disrupted due to technical glitches, market liquidity significantly worsens. Our findings suggest that the consolidated feed continues to be a crucial component of today's market data infrastructure, shedding light on market data regulation.

1 Introduction

Trading has become increasingly fast and fragmented in the past two decades. Take the U.S. equities market as an example. Trades are executed at sub-millisecond frequencies[1] across 16 public exchanges, dozens of broker-dealer internalizers, and dark pools. In such a trading environment, having access to low-latency market data is crucial for various market participants to implement their trading strategies: Market makers need it to re-price their quotes on time to avoid adverse selection; Institutional investors to find the best liquidity to execute their trades and arbitragers to exploit short-lived arbitrage opportunities.

Two types of market data exist in the U.S. equities market: the consolidated feed and direct feeds. While the former is mandated by regulation, the latter are expensive, proprietary products of exchanges and thus have lower latency and contain more information.[2] Market participants argue that the discrepancy between the two feeds has created an unfair "two-tiered" market with the haves (direct feeds subscribers such as high-frequency traders, HFTs) and have-nots (the consolidated feed subscribers). To address such a concern, the U.S. Securities and Exchange Commission (SEC) proposed a new rule named Reg NMS II in February 2021, aiming to improve the consolidated feed and level the playing field of market data. However, the rule has caused a fierce backlash from exchanges.[3] Across the Atlantic, while the European-wide consolidated feed is not yet in place, the European Securities and Markets Authority (ESMA) is actively laying the groundwork for its introduction. However, details about the feed, concerning its content and speed, are under significant discussion.

Objective. Exchanges argue that since the consolidated feed's main users are "display users" or human traders, further speed enhancements at sub-second frequencies might not benefit them.

[1]For example, Menkveld (2018) analyzes a sample of Nasdaq trades in October 2010 and finds that twenty percent of trades arrive in sub-millisecond clusters.

[2]See Section 3 below for details about the speed and content difference between the consolidated feed and direct feeds.

[3]SEC approved the rule in December 2020 but was sued by Nasdaq, NYSE, and Cboe. Thus the implementation of the rule is now blocked. See "Nasdaq, NYSE Sue SEC to Block Market Data Overhaul.", *The Wall Street Journal*, February 9, 2021.

As for latency-sensitive traders, they will always opt for faster direct feeds and be hardly affected either.[4] However, anecdotal evidence shows that as the consolidated feed becomes faster, it is more used by algorithmic traders.[5] Moreover, regulatory filings show that about 45% dark pools do not use any direct feeds but the consolidated feed instead.[6] Even for low-latency traders, they typically subscribe to the consolidated feed for data integrity check (CFTC and SEC 2010).[7] Last, off-exchange trade reports from the consolidated feed might contain useful information even though they appear with a significant delay (Ernst, Sokobin, and Spatt 2021). The key questions are: Does the consolidated feed matter in today's market? How does a faster consolidated feed affect market liquidity and trading? Our paper aims to shed light on the above questions and contribute to the ongoing regulatory debates.

Approach. We exploit an exogenous technological shock to the consolidated feed on October 24, 2016, which significantly cut its processing latency: The median latency drops from about 350 microseconds to less than 20 microseconds. The U.S. equities market features two consolidated feed streams: one for Nasdaq-listed stocks (Tape C) managed by the Nasdaq Security Information Processor ("Nasdaq-SIP"), and another for NYSE-listed and other regional exchange stocks (Tapes A and B) managed by the "NYSE-SIP". The latency reduction only happened at the Nasdaq-SIP and thus allows us to perform a clean difference-in-difference (DiD) analysis based on a matched sample of Nasdaq-listed and NYSE-listed stocks. Should the consolidated feed's speed matter for market participants, we would expect to observe liquidity and trading changes for Nasdaq-listed stocks post-upgrade, compared to NYSE-listed stocks.

In addition, we utilize the unique geography of the consolidated feed for further identification. The Nasdaq exchanges and Nasdaq-SIP are housed in Carteret, New Jersey, while the NYSE exchanges and NYSE-SIP are roughly 35 miles away in Mahwah, New Jersey. Consequently, for Nasdaq-listed stocks, the traveling latency of a message from the Nasdaq exchanges to Nasdaq-SIP

[4]"Building the SIP Autobahn", *Nasdaq*, May 20, 2021.

[5]"Consolidated Market Data Feeds Gain Traction in Algo Trading and Fixed Income", *Finextra*, January 2019.

[6]See, for example, "Dispelling the Complementary Product Theory for Market Data", *Nasdaq*, August 8, 2020.

[7]Aldrich, Grundfest, and Laughlin (2017) argue that during the 2010 Flash Crash, disruptions to the consolidated feed might cause traders to withdraw from the market.

is small, with the SIP's processing time being the predominant component of total latency. With the upgrade slashing the processing time by over 90%, the total latency from the Nasdaq exchanges sees a more substantial relative decrease compared to that from NYSE exchanges.[8] Hence, we expect a more noticeable DiD impact on Nasdaq's own exchanges.

Additionally, we analyze instances where technical glitches disrupted the consolidated feed. Since all glitches only affected one stream of the consolidated feed at a time, we can employ the same DiD approach as before. To assess market liquidity during the glitches, we use direct feed data from MayStreet, which is necessary because the consolidated feeds, like the NYSE TAQ database, are unavailable.

Findings. Surprisingly, we find that a faster consolidated feed does not improve but slightly deteriorates market liquidity. Following the Nasdaq-SIP upgrade, which considerably decreases the latency of the consolidated feed for Nasdaq-listed stocks, our analysis indicates that compared to NYSE-listed stocks, there's a relative increase in the quoted spread, effective spread, and price impact for Nasdaq-listed stocks. To quantify, across approximately 600 matched stocks, the DiD estimates show that the three liquidity measures rose modestly by about 6.32%, 4.76%, and 5.53%, respectively. Additionally, the dampening effect on liquidity seems to vary by exchange, with the quoted spread on the Nasdaq exchange showing a larger increase than that on the NYSE, aligning with our initial conjectures.

Why would a faster consolidated feed lead to worse market liquidity? We find that the upgrade results in an increase in high-frequency quoting activities as evidenced by significant rises in order-to-trade and cancel-to-trade ratios—nearly 12.99% and 16.11% respectively based on the DiD estimates. The results suggest that HFTs may now face more competition from slow traders (e.g., buy-side execution algorithms) as the speed edge of direct feeds over the consolidated feed greatly vanishes. Consequently, HFT market makers have to re-price their quotes more frequently and quote wider spreads to compensate for larger adverse selection costs, reflected in increased price impacts.

[8] See Figure 1 for a visual illustration and Section 5.2 for a detailed explanation.

To provide more direct evidence, we examine two trading inefficiencies slow traders might face using the consolidated feed versus direct feeds. The first metric is the number of sub-optimally routed trades—those executed on maker-taker exchanges while better prices are available on inverted exchanges. Post-upgrade, there is a significant reduction in the number of such trades for Nasdaq stocks relative to NYSE stocks. The second metric is the time it takes for traders to take liquidity on inverted exchanges. Specifically, we calculate the fill time of liquidity-improving orders on inverted exchanges and find that the differences between NYSE and Nasdaq stocks are similar in the low-latency regions (below 1 millisecond) but start to expand as we move to the higher-latency regions (around and above 1 millisecond) after the upgrade. It suggests that while a faster consolidated feed does not affect the fastest traders like HFTs, it benefits slower traders, allowing them to access inverted liquidity quicker.

When the consolidated feed is disrupted due to technical glitches, we document that market liquidity deteriorates significantly, especially in terms of market trading volume and order-book depth. Specifically, averaging across a sample of three glitch events, the trading volume of affected stocks falls by 17.81% more than unaffected stocks. Similarly, depth at best prices and cumulative depth across the five best price levels of the affected stocks drop by about 21.20% and 13.80% more for treatment stocks than control stocks respectively.

Our paper highlights the significant role of the consolidated feed in a fast and fragmented market. In particular, speed improvement to the feed may cause an unintended and adverse effect on market liquidity as it alters the competition between traders at different speeds. Therefore, any future proposals for changes in speed should be approached with caution.[9] While our analysis is focused on the US equities market, it can generalize to other markets with similar characteristics. For example, in the EU capital market, where trading venues are more geographically dispersed, latency in data feeds could have an even greater influence on market liquidity and trading activities.

Outline. The paper proceeds as follows. In Section 2, we discuss the related literature and our

[9] It's worth mentioning that our analysis cannot directly address the impact of the SEC's new data overhaul plan, which involves adding more information such as depth to the consolidated feed or introducing competing consolidators.

contribution. In Section 3, we introduce institutional details about the unique structure of market data in the U.S. equities market. We describe our dataset and define variables in Section 4. In Section 5, we detail our identification strategy used for the empirical analysis. We discuss the results in Section 6 and conclude in Section 7.

2 Literature

Our paper relates to several strands of literature. First, it is directly related to the literature studying the effect of market data on trading and market quality. Brogaard and Brugler (2020) examine events where exchanges start to charge a fee for their direct feeds for the first time. They find that the introduction of data fees leads to a significant fall in the market volume of the fee-charging exchange and it is mainly due to it having less time at the National Best Bid and Offer (NBBO) and getting fewer inter-market sweep orders. Hendershott, Rysman, and Schwabe (2020) study the event when NYSE introduced its new data product, NYSE Integrated Feed, and show that there is a complementary relationship between the exchange's proprietary data sales and trading activity: Firms increased their share of trading on NYSE after the introduction. Our paper differs from theirs in that we focus on the consolidated feed instead of the exchanges' direct feeds. By examining exogenous event affecting its speed, we show that the consolidated feed matters in today's market and play a significant role in shaping trading behavior and market liquidity, contributing to the ongoing regulatory debate over market data reform. One exception is Ye, Yao, and Gai (2012) where the authors examine an older speed upgrade to the Nasdaq-SIP and find it has no impact on overall market liquidity. However, they do not study its impact on trading inefficiencies.

Second, our paper adds to the literature on the difference between the consolidated feed and direct feeds. O'Hara, Yao, and Ye (2014) show that odd-lot trades were missing in the consolidated feed[10] which account for a large share of trading volume and are informed. They conjecture that odd-lot traders are used by the informed traders to hide their trading intention from the consol-

[10]Odd-lot trades were added into the consolidated feed after a SIP reform in 2013.

idated feed. Battalio, Corwin, and Jennings (2016) analyze a sample of high-priced stocks and show that the exclusion of odd-lot orders from the consolidated feed results in some trades being filled at worse prices. Ding, Hanna, and Hendershott (2014) compare NBBO constructed from the consolidated feed and NBBO constructed by adding direct feeds from Nasdaq, BATS, and Direct Edge exchanges. They find that the dislocation between the two NBBOs can happen quite frequently for active stocks but its duration is quite short. So the cost for low-frequent traders is small. Similarly, Bartlett and McCrary (2019) use the exchange timestamp included in the consolidated feed to construct NBBO with zero-latency and show that profit from direct feed arbitrage is not economically significant. Hasbrouck (2021) compares the price discovery contribution of SIP compared to theoretically constructed direct feeds and finds the latter dominates at high frequency. Bartlett, McCrary, and O'Hara (2023) show that odd-lot quotes included in the direct feeds but not the consolidated feed are predictive of future prices. Our paper contributes to the existing studies by examining the event when the speed advantage of direct feeds over the consolidated feed largely narrows. Compared with the approach used in the existing studies that statically compares the two feeds, our paper looks at a real-life event so that it can incorporate the effect of traders' equilibrium responses to a faster consolidated feed.

Third, our paper closely relates to studies that examine the economics of trading speed[11]. Based on the adverse selection channel, theoretical models have shown that the impact of trading speed depends on who becomes faster. Market liquidity worsens if short-term informed arbitrageurs become faster (Biais, Foucault, and Moinas 2015; Budish, Cramton, and Shim 2015; Foucault, Hombert, and Roşu 2016), improves if market makers become faster (Hoffmann 2014; Jovanovic and Menkveld 2016), and depends on market conditions (e.g., news arrival frequency and the presence of noise traders) when both trader groups become faster (Menkveld and Zoican 2017). Empirical studies largely support the theoretical predictions above. For example, Brogaard, Hagströmer, Nordén, and Riordan (2015) show that market liquidity improves when market-makers opt for a speed upgrade of their co-location server to the exchange. Shkilko and Sokolov (2020) find that

[11]See Menkveld (2016) for a comprehensive review.

market liquidity improves when the microwave network between Chicago and New York is disrupted by weather conditions, which makes high-frequency arbitrageurs slower. Our paper adds to the literature by examining the exogenous shock where a faster consolidated feed increases the trading speed of slow traders (e.g., buy-side execution algorithms). We find the speed upgrade has a negative impact on overall market liquidity and show that it is likely due to more intensive competition between HFTs and slower informed traders as evidenced by higher price impacts, more quoting activities, and better routing decisions.

3 Institutional Background

Before diving into our data and methods, we first introduce the institutional framework of the U.S. equities market data structure, which is important for understanding our identification approach.

3.1 A Two-Tiered Market Data Structure

Broadly speaking, there are two types of market data in today's U.S. equities market. First, there is the consolidated feed mandated by the SEC. The consolidated feed is disseminated by the Security Information Processors (SIPs), which, on a near real-time basis, collect trades and top-of-book quote updates from all national securities exchanges such as NYSE and Nasdaq, aggregate them into consolidated tapes and disseminate them to their subscribers. In addition to trades and quote updates, the SIPs disseminate critical regulatory information including the National Best Bid and Offer (NBBO), Limit Up-Limit Down (LULD) price bands, short sale restrictions, and halts. Subscribers of the SIPs form a heterogeneous group, including TV and media, broker-dealers, investment advisors, and algorithmic traders, and they are charged varying rates depending on their specific type. Revenue of the SIPs is then shared across all participating exchanges based on their contribution to market trading volume and depth at the NBBO.[12]

[12] See "SIP Accounting 101", Nasdaq, March 25, 2021, at https://www.nasdaq.com/articles/sip-accounting-101-2021-03-25

Second, there are direct feeds sold by exchanges as their proprietary data products. Compared with the consolidated feed, direct feeds are faster and contain more information such as depth-of-book information, odd-lot quotes, and auction imbalance information. As a result, sophisticated traders such as high-frequency traders use direct feeds as key inputs to their trading algorithms. However, direct feeds are prohibitively expensive for unsophisticated traders. A 2019 report published by IEX on market data cost estimates that the total annual subscription fee to all national securities exchanges' direct feeds sum up to around 1.15 million US dollars, excluding costs for physical and logical connectivity needed to receive the feeds (IEX 2019).

Such a two-tiered market data structure has spurred heated debates over whether it is unfair to small market participants as they in a way suffer from information asymmetry due to inferior market data. The SEC has been improving the consolidated feed to bring them closer to direct feeds in terms of speed and content. Besides the processing latency reduction, which is the focus of the paper, the most recent proposal by SEC called NMS 2.0 plans to add depth-of-book information (up to five levels), odd-lot quotes, and auction imbalance information to consolidated feeds.[13]

3.2 Unique Features of the Consolidated Feed

There are two unique features of the consolidated feed for the U.S. equities market. First, the feed is not disseminated by one SIP, but by two separate ones. Specifically, one SIP is managed by NYSE ("NYSE-SIP") and responsible for disseminating the consolidated feed for securities listed on NYSE (Tape A Securities) or other regional exchanges and their successors (Tape B Securities, e.g., ETPs listed on NYSE ARCA). Instead, the other SIP is managed by Nasdaq ("Nasdaq-SIP") and responsible for disseminating the consolidated feed for securities listed on Nasdaq (Tape C Securities). The bottom line is that there is a separation between where security is traded and where the trade is reported: Trades and quotes of stock from any national exchange have to be reported to the NYSE-SIP as long as its listing exchange is NYSE, and to the Nasdaq-SIP if its listing exchange is Nasdaq.

[13] See details of the proposed rule at https://www.sec.gov/news/press-release/2020-311

Figure 1. The New Jersey "Equity Triangle". This figure sketches the geographical locations of major US equity exchanges (in rectangles) and two SIPs (in circles). Besides, it shows the estimated traveling latencies between three data centers and processing latencies at the two SIPs before the upgrade event on October 24, 2016. Latency numbers are taken from "Time is Relativity: What Physics Has to Say About Market Infrastructure", *Nasdaq*, April 9, 2021.

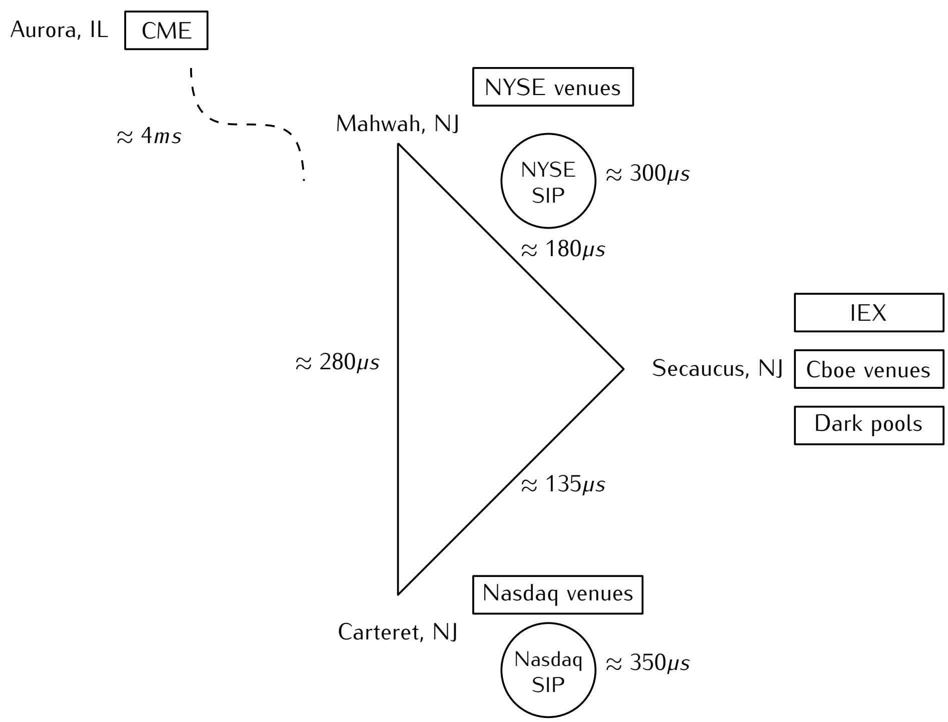

Another feature of the feed is its unique geography. As depicted in Figure 1, all major US equities exchanges and the two SIPs are located in a triangular area in New Jersey, with (1) Nasdaq venues (Nasdaq, BX, and PSX) and the Nasdaq-SIP at Carteret, NJ, (2) NYSE venues (NYSE, NYSE Arca, NYSE American) and the NYSE-SIP in Mahwah and (3) Cboe venues (BZX, BYX, EDGX and EDGA) in Secaucus. The geography of SIPs and exchanges creates what's called the "geographical latency" of the consolidated feed. For example, a trade or quote update for a Nasdaq stock from NYSE venues will have to first travel from Mahwah to Carteret to get processed by the Nasdaq-SIP.

3.3 Details about the SIP Upgrade Event

The two SIPs have received several technological upgrades over the recent years, resulting in significant reductions in their processing latency. As Figure 2 shows, there were three major up-

Figure 2. SIP Speed Upgrades. Starting from August 2015, each TAQ message has two timestamps: Participant Timestamp (when the message is registered at the exchange from which it originates) and SIP Timestamp (when the message is disseminated by the SIP). Thus we can compute the SIP latency by the difference between the two timestamps. To compute the processing latency of the NYSE-SIP, we use quotes in General Electric originating from the NYSE exchange so that there is minimal traveling latency. By the same token, to compute the processing latency of the Nasdaq-SIP, we use quotes in Apple originating from the Nasdaq exchange. The figure plots the daily median latency for NYSE-SIP and Nasdaq-SIP respectively.

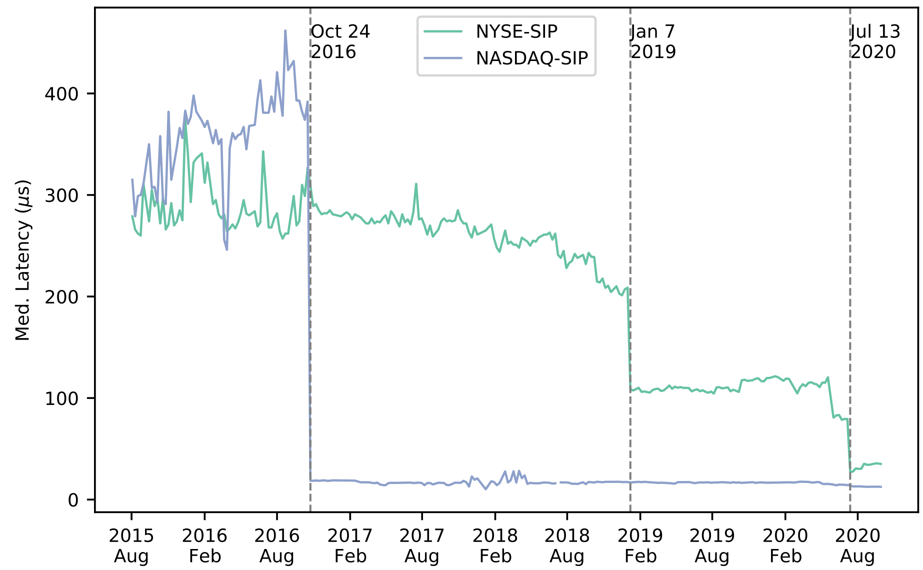

grades: (1) Nasdaq-SIP on October 24, 2016, (2) NYSE-SIP at the end of 2018 and, (3) NYSE-SIP on July 13, 2020. In our analysis below we focus on the Nasdaq-SIP upgrade on October 24, 2016, as the latency reduction is most significant: Its median processing latency drops more than 90% from about 350 microseconds to less than 20 microseconds. Specifically, Nasdaq implements the technology upgrade of its SIP by migrating it to the Nasdaq Financial Framework and INET, the proprietary technology behind Nasdaq exchanges. As a result, not only was there a significant decrease in the Nasdaq-SIP's processing latency, but also an increase in its resiliency and reliability, capacity, scalability, and message efficiencies.[14]

Table 1. Recent Market-Wide SIP Glitches. This table lists some key information about the three market-wide SIP glitch events that happened in recent years.

Date	Start and End Time	Duration	SIP	Market-Wide Trading Halt
January 3, 2013	13:33 - 13:51	18 minutes	Nasdaq-SIP	No[a]
October 30, 2014	13:07 - 13:34	27 minutes[b]	NYSE-SIP	No[c]
August 12, 2019	15:15 - 15:27	12 minutes[b]	NYSE-SIP	No

[a] There is no market-wide trading halt. EDGX and EDGA halted trading for Nasdaq-listed stocks after 13:42.

[b] In both events, The NYSE shifted operations to its disaster recovery site in Chicago after the glitch was solved.

[c] Some dark pools, including ITG Posit and Goldman Sachs' Sigma X, which uses the NYSE SIP were closed during the glitch period.

3.4 Details about SIP Glitch Events

While SIP glitches are not rare, there are only a few with market-wide impact. Going through all system alerts published by the two SIPs[15], we find three market-wide SIP glitch events in recent years: (1) the Nasdaq-SIP glitch on January 3, 2013, (2) the NYSE-SIP glitch on October 30, 2014, and (3) the NYSE-SIP glitch on August 12, 2019. In Table 1, we provide some basic information about the three SIP glitches such as the start and end time, duration, and whether there was a trading halt during the glitch period. We offer a succinct account for each SIP glitch event in Appendix B.

We would like to mention that the three SIP glitch events share two important features. First, unlike the Nasdaq "Flash Freeze" event on August 22, 2013, which resulted in a three-hour trading halt for Nasdaq-listed stocks, there was no market-wide trading halt during any of the three SIP glitch events. Second, while the consolidated feed was disrupted during glitches, direct feeds from exchanges operated normally. Although the duration of the three SIP glitch events is fairly short and their economic impact might be small, they provide us an opportunity to empirically study what happens to the market when an important piece of data infrastructure is missing and there is maximum information asymmetry: Sophisticated traders such as HFTs who do not rely on the consolidated feed but subscribe to direct feeds from exchanges are not directly affected; In

[14] See "Nasdaq Sets the Record Straight About the SIP", Nasdaq, October 25, 2016, https://www.nasdaq.com/articles/nasdaq-sets-record-straight-about-sip-2016-10-25.

[15] Market data alerts for the NYSE-SIP and Nasdaq-SIP can be found at https://www.ctaplan.com/alerts# and https://www.utpplan.com/vendor_alerts respectively.

contrast, retail brokers and dark pools which rely on the consolidated feed will receive unreliable or no pricing information at all.

4 Data and Variables

In the following section, we first introduce the datasets used in our analysis of both the SIP upgrade and glitch events. Then, we define our market liquidity and trading measures for each event type.

4.1 Data Sources

Datasets used for analyzing SIP upgrade events The first dataset we use is the NYSE TAQ database[16], which is essentially NYSE's version of the SIP feeds and includes trades and top-of-book quote updates from all national securities exchanges. With the NYSE TAQ, we calculate common liquidity and trading variables such as bid-ask spread, top-of-book depth, and trading volume. In addition, we use the SEC's MIDAS dataset, which collects direct feeds from all national securities exchanges and publishes several useful aggregated trading variables.[17] For example, for each stock and exchange combination, MIDAS reports the daily count and volume of new order submissions, count of cancel messages, hidden volume, and odd-lot volume.

Our sample covers a matched sample of 296 Nasdaq-listed stocks and the same number of NYSE-listed stocks over the period from August 29, 2016, to December 16, 2016. So the sample period spans a four-month window, two months before and two months after the event date of October 24, 2016. A four-month window length is chosen to strike a balance between a too-long window which might include other events and a too-short window which might not generate sufficient statistical power. As a robustness check, in Appendix A.3 of the appendix, we shorten the window to two months, one month before and one month after the event date of October 24, 2016, and the results do not change significantly.

[16]The NYSE TAQ dataset is accessed through WRDS (Wharton Research Data Services).
[17]https://www.sec.gov/marketstructure/downloads.html

Datasets used for analyzing SIP glitch events When analyzing SIP glitch events, we use instead direct feeds data collected by MayStreet, a US data company and supplier of SEC's MIDAS. Direct feeds from different exchanges have different formats.[18] For some direct feeds such as NYSE OpenBookUltra, they contain only level book messages which represent an update on a single level of the order book. In other words, we can not see individual order submissions and cancellations as all orders at a particular price are communicated in a consolidated format. In contrast, direct feeds from other exchanges contain full order-book event messages, i.e., new order submissions, cancellations, and executions. In either case, we can use the direct feed messages from each exchange to build the exchange-specific limit order book.

We obtain the "consolidated" limit order book for the whole market by aggregating limit order books of all exchanges. We use the consolidated limit order book to compute several market-wide and high-frequency liquidity and trading measures. It is worth noting that direct feeds from exchanges are necessary in order to study the SIP glitch events as, by definition, SIP feeds (such as the NYSE TAQ) are either unreliable or unavailable during such events. As with the analysis of the Nasdaq-SIP upgrade, we include a matched sample of Nasdaq-listed stocks and NYSE-listed stocks. In addition, for each SIP glitch event, we include observations in the glitch period and a half-an-hour window before the start of the glitch.

4.2 Liquidity and Trading Variables

Constructing variables for analyzing SIP upgrade events We calculate the following measures capturing different aspects of market liquidity.

[18]Specifically, we use the following direct feeds: BATS BZX Multicast Pitch, BATS BYZ Multicast Pitch, CHX Book Feed, DirectEdge EDGA Multicast EdgeBook Depth, DirectEdge EDGX EDGX Multicast EdgeBook Depth, Nasdaq BX TotalView-ITCH, Nasdaq TotalView-ITCH, Nasdaq PSX TotalView-ITCH, NYSE MKT OpenBook Ultra, NYSE ARCA ARCABook, ARCA Trades, NYSE OpenBook Ultra, NYSE Trades, National Stock Exchange Multicast Depth of Book.

- *RQS*, relative quoted spread, is defined as:

$$RQS_t = \frac{NBO_t - NBB_t}{Mid_t}, \tag{1}$$

where t is the timestamp of the current snapshot of the order book. NBO_t and NBB_t are the national best offer and bid respectively. Mid_t is the midquote, which is simply $(NBO_t + NBB_t)/2$. *RQS* measures the cost of a round trip of small trades. However, *RQS* can be a poor proxy for the actual transaction cost traders pay for two reasons. First, for a large trade that "walks up the order book", i.e., is executed against quotes at multiple price levels, its average transaction price is worse than the prevailing NBBO. Second, it is common for trades to be executed at prices better than NBBO: Off-exchange trades can either be executed at the NBBO midquote in dark pools or receive price improvement from wholesalers like Citadel and Virtu; On-exchange trades, instead, can be executed against hidden orders priced better than NBBO.

- *RES*, relative effective spread, is defined below:[19]

$$RES_t = \frac{d_t\,(p_t - Mid_t)}{Mid_t}, \tag{2}$$

where t indexes trades. d_t is the trade direction indicator. p_t is the actual transaction price for trade t.[20] Mid_t is the prevailing midquote *just before* the trade. *RES* captures the actual transaction cost traders pay.

[19]When computing trade-based liquidity measures such as effective spread, we drop trades with Stock-Option Trade, Average Price Trade, Derivatively Priced Trade, etc. These trades often have execution prices far off from the prevailing market price, skewing their trade-based liquidity measures. Thus we drop those trades when computing the daily metrics. Specifically, we only keep trade records with trade conditions "@", "F", "I", "F I", for Tape A securities and"@F", "I", "@F I", "@ I", for Tape C securities. Moreover, when deciding the trade direction, we use the Exchange Timestamp instead of the SIP Timestamp.

[20]For trades with multiple executions, we use the volume weighted average price.

- *RRS*, relative realized spread, is defined below:

$$RRS_t = \frac{d_t\,(p_t - Mid_{t+\Delta_t})}{Mid_t}, \tag{3}$$

  where $Mid_{t+\Delta_t}$ is the prevailing midquote after a time interval of Δ_t after the trade. We pick the common choice of 30 seconds. The relative realized spread is essentially the relative effective spread less the price impact and is a crude proxy for the profits of market makers[21].

- In addition to the liquidity measures above, we compute two other measures capturing the quantity aspect of liquidity: dollar depth at NBBO (*Depth*), and dollar trading volume (*Vlm*). A deeper dollar depth at NBBO makes large trades cheaper. While trading volume depends on various factors including volatility and information. However, given the same volatility and information level, better market liquidity leads to a larger trading volume as traders find it cheaper to trade and realize their private values.

In addition to market liquidity measures, we calculate several variables that capture the trading behavior of market participants.

- *ISOShr*: trading volume via inter-market sweep order (ISO) as a fraction of total trading volume. The SEC mandates a no "trade-through" rule, meaning that an incoming marketable order has to be executed against the best prices across all exchanges. For example, if the prevailing best ask is 10.02 and 10.01 at Nasdaq and NYSE respectively, a market buy order routed to Nasdaq can not be executed as there is a better price at NYSE. In such a case, Nasdaq has the obligation to re-route the order to NYSE. Although the rule guarantees the best prices, it is not desirable for traders who wish to execute a large quantity immediately. An ISO order is used by such traders to sweep liquidity at other exchanges beyond the NBBO.[22]

[21]The relative realized spread is only a crude proxy for market maker profit. For example, it does not include the rebates market makers receive from the exchanges and various costs from co-location, exchange data subscription, and fixed IT costs.

[22]ISO trades can be identified in the TAQ with a trade indicator of "F".

- *OddlotShr*: odd-lot trading volume as a fraction of total trading volume. Odd-lot trades refer to transactions involving fewer than the standard "round lot" of 100 shares. While commonly attributed to retail investors, they can also stem from high-frequency trading strategies.[23]

- *Cancel/Trade* and *Order/Trade*: the ratio of cancel order count to trade count and the ratio of the volume of add limit orders to trading volume respectively. While the former captures the cancellation activities, the latter measures new quote submissions. Hagströmer and Nordén (2013) show that market-making HFTs have substantially higher quote-to-trade ratios than other fast traders as they constantly reprice their quotes in order to avoid adverse selection. So an increase in the quote-to-trade ratio is likely to reflect an increase in market-making-related activities.

We'd like to stress one important technical detail. Since August 2015, SIP trade and quote messages have started to include two timestamps: Participant Timestamp, the time when the message is registered at the originating exchange, and SIP Timestamp, the time when the message is disseminated by the SIP. As argued above in Section 3 on institutional details, SIP trade and quote messages can be subject to delays due to either processing latency at the SIP or traveling latency from the originating exchange to the SIP. Thus the SIP Timestamp will be later than the corresponding Participant Timestamp and the NBBO constructed based on the SIP Timestamp ("SIP-NBBO") will lag that based on the Participant Timestamp ("Participant-NBBO").

The discrepancy between the SIP-NBBO and Participant-NBBO can lead to biases in certain liquidity measures if they are computed based on the SIP-NBBO included in the NYSE TAQ database. For example, trade-related liquidity measures such as effective spread require one to first infer the direction of a trade, i.e., whether it is buyer- or seller-initiated. One popular trade classification algorithm is Lee and Ready (1991): Trades with a transaction price higher (lower) than the prevailing midquote will be classified as buy (sell) trades. So latency embedded in the SIP-NBBO can result in misalignment between trades and their actual prevailing NBBOs and thus wrong classifications.

[23] Odd-lot trades can be identified in the TAQ with a trade indicator of "I".

Table 2. Summary Statistics: The Nasdaq-SIP Upgrade Event on October 24. 2016. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact, and relative realized spread respectively, and are all in basis point. *Depth* is NBBO depth in thousand dollars. *Vlm* is trading volume in million dollars. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *ISOShare* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. The sample period is from August 29 to December 16, 2016.

Variable	N	Mean	SD	Min	50%	Max
RQS (bp)	42900	8.55	9.01	0.85	6.19	310.27
RES (bp)	42900	2.65	3.15	0.28	1.85	204.06
RPI (bp)	42900	2.35	2.08	-12.02	1.83	100.72
RRS (bp)	42900	0.29	2.12	-44.76	0.01	128.58
Depth ($ thousand)	42900	160.11	505.77	6.94	76.59	14241.81
Vlm ($ million)	42900	104.23	238.14	0.20	46.24	11120.27
ISOShr (%)	42900	35.03	7.23	3.66	34.96	89.36
OddlotShr (%)	42900	10.75	6.72	0.06	9.65	56.24
Cancel/Trade	42900	23.79	11.98	3.84	21.14	252.44
Order/Trade	42900	38.04	23.76	5.48	32.31	443.62

Another measure that might be biased based on SIP-NBBO is one exchange's NBBO depth contribution. Imagine there are two exchanges, Exchange A and Exchange B, and the former always sets the new NBBO first and the latter follows. However, due to geographical latency, Exchange B's quote updates might be processed by the SIP first before Exchange A's arrive. Thus Exchange B will always be the exchange that sets the new NBBO. To make things more complicated, SIP processing time is significantly reduced in the SIP upgrade event we study, which might mechanically change the liquidity measures. Given the foregoing reasons, we construct the BBOs for all exchanges and then the NBBO based on the Participant Timestamp. In addition, when matching trades with their prevailing NBBOs, we use trades' Participant Timestamp as well.

All liquidity measures are first computed at the tick-by-tick frequency and later aggregated to a lower frequency (e.g., daily for the SIP upgrade events). We aggregate stock variables such as *RQS* and *Depth* by computing their time-weighted averages and flow variables such as *RES* by their dollar-volume weighted averages. Table 2 reports the summary statistics of the above liquidity and trading metrics for the matched sample of Nasdaq-listed and NYSE-listed stocks around the Nasdaq-SIP upgrade event.

Table 3. Summary Statistics: SIP Glitch Events. This table reports the summary statistics for all liquidity variables used in the regression for pooling all three sip glitches. *RQS*, *RES* and *RRS* stands for relative quoted spread, effective spread, and realized spread respectively, and are all in basis point. *Vlm* is dollar volume in thousand dollars. *DepthNBBO* is NBBO depth in thousand dollars. *Depth5Lvl* is cumulative depth across the five best price levels of the order book. The summary statistics are computed based on time-series averages over the sample period.

	N	Mean	SD	Min	50%	Max
RQS (bp)	235564	13.80	18.45	0.52	9.08	730.64
RES (bp)	129309	3.32	4.81	0.00	2.04	319.86
RRS (bp)	129309	1.38	6.62	-366.81	0.94	323.56
Vlm ($ million)	235564	33.30	201.72	0.00	1.82	45054.15
DepthNBBO ($ thousand)	235564	84.50	347.56	0.07	32.52	34608.91
Depth5Lvl ($ thousand)	235564	470.26	1387.41	4.25	149.43	42915.56

Constructing variables for analyzing SIP glitch events For the analysis of SIP glitch events, we focus on their impact on market-wide liquidity and compute several common liquidity measures. Specifically, we compute the relative quoted spread (*RQS*), relative effective spread (*RES*), and relative realized spread (*RRS*). In addition, we compute the dollar trading volume (*Vlm*), top-of-book depth (*DepthNBBO*), and cumulative depth across five best prices of the order book (*Depth5Lvl*). Note that as we obtain direct feeds from all exchanges during the SIP glitches, we can not only compute the top-of-book depth but also the depth across several levels of the order book. Table 3 reports the summary statistics for all liquidity variables used in the regression for pooling all three SIP glitches.

5 Identification Strategy

In the following section, we detail our identification strategy. Specifically, the unique features of the consolidated feed for the U.S. equities market allow us to construct a matched sample of NYSE-listed stocks and Nasdaq-listed stocks and perform difference-in-difference estimations.

5.1 Nasdaq-Listed Stocks vs. NYSE-Listed Stocks

Due to the technological upgrade of the Nasdaq-SIP on October 24, 2016, its processing latency decreased significantly while NYSE-SIP's barely changed (See Figure 2). As the Nasdaq-SIP disseminates the consolidated feed for Nasdaq-listed stocks while the NYSE-SIP for NYSE-listed stocks, the Nasdaq-SIP upgrade event allows for a clean difference-in-difference (DiD) analysis: After the Nasdaq-SIP upgrade, the consolidated feed becomes much faster for Nasdaq-listed stocks but not for NYSE-listed stocks. If a faster consolidated feed has any impact on liquidity or trading, we should observe our proxies changes for Nasdaq-listed stocks *relative to* NYSE-listed stocks.

To implement the DiD analysis, we construct a matched sample of Nasdaq-listed and NYSE-listed stocks. We follow Brogaard and Brugler (2020) and include all common stocks[24] listed on either NYSE or Nasdaq and exclude stocks with dual-class shares and a market capitalization below $500 million. More importantly, we exclude stocks involved in the SEC's Tick Size Pilot Program to avoid its confounding effect. The program started in October 2016 and was conducted by the SEC to assess the impact of wider tick sizes on the liquidity and trading of certain small-capitalization companies ("Pilot Securities"). Pilot Securities are divided into one control group and three test groups. While tick sizes of stocks in the test group remain at $0.01, those in the test groups increase from $0.01 to $0.05 either for their trading or quoting or both.[25] After excluding all Pilot Securities, including the control group, we are left with 296 Nasdaq-listed stocks and 633 NYSE-listed stocks.

We match the 296 Nasdaq-listed stocks with the same number of NYSE-listed stocks on price, trading volume, market capitalization, and industry. The first three matching variables are the daily averages during the month before the event period, that is, between August 1, 2016 and August 26, 2016. In addition, as companies in certain industries have a preference for listing either on Nasdaq or NYSE (e.g., technology companies are more likely to list on Nasdaq), we follow Brogaard and Brugler (2020) and add Fama and French 12 industry classification as a further

[24] Common stocks have a CRSP share code of 10 or 11.
[25] We refer readers to the program's official website for details (https://www.sec.gov/ticksizepilot)

Table 4. Propensity Score Matching: SIP Speed Upgrade. This table reports results from the propensity score matching for the SIP speed upgrade. The treatment group consists of 296 Nasdaq-listed stocks that are matched with 296 NYSE-listed stocks on price, trading volume, market capitalization, and Fama and French 12 industry classification. We use one-to-one nearest neighbor propensity score matching (PSM), without replacement.

Variable	Sample	N	Mean	SD	10%	25%	50%	75%	90%
Price ($)	Control	296	62.03	54.25	14.67	25.52	47.14	79.53	130.23
	Treatment	296	67.35	110.38	10.01	22.08	45.14	79.68	119.85
MarketCap ($ billion)	Control	296	15.32	26.43	1.93	3.35	6.07	15.24	33.27
	Treatment	296	16.95	51.23	1.28	3.08	4.93	11.57	30.06
DollarVolume ($ million)	Control	296	109.34	160.25	15.92	28.75	64.12	124.99	229.15
	Treatment	296	114.08	241.66	13.74	24.99	48.36	102.09	242.71
PSM Score	Control	296	0.41	0.15	0.25	0.29	0.46	0.51	0.61
	Treatment	296	0.44	0.16	0.25	0.30	0.48	0.55	0.62

matching variable.[26] Moreover, we use one-to-one nearest neighbor propensity score matching (PSM), without replacement. Table 4 reports the matching results and shows that matching is successful as all three matching variables and the propensity score of the two samples have quite similar support.

The DiD identification approach requires the standard parallel trends assumption. Figure A1 in Appendix A.1 plots the time series of the liquidity and trading metrics around the Nasdaq-SIP upgrade event. Visual inspection shows that all variables for the treatment group (i.e., Nasdaq-listed stocks) evolved in a similar fashion to the control group (i.e., NYSE-listed stocks) before the event date, suggesting that the parallel trend assumption is supported. In Appendix A.2, we further perform a placebo test with a two-month window before the event date to provide statistical evidence.

Based on the matched sample described above, we estimate the difference-in-difference (DiD) regressions as follows:

$$metric_{i,t} = \alpha_i + \beta After_t + \gamma After_t \times NasdaqStock_i + \epsilon_{i,t} \tag{4}$$

[26] Perhaps it is worth noting that they use Fama and French 48 industry classification. The reason for me to use the simpler version is due to the relatively small size of the sample. Not every industry has stocks from both tapes in our sample, which makes it impossible for the logit regression to converge.

Table 5. Propensity Score Matching: SIP Glitch Events This table reports results from the propensity score matching for the three SIP glitch events. For each event, the treatment group consists of 1200 randomly chosen stocks whose consolidated feeds are affected by the SIP glitch and are matched with 1200 unaffected stocks on price, trading volume, market capitalization and Fama and French 12 industry classification. We use one-to-one nearest neighbor propensity score matching (PSM), without replacement.

Variable	Sample	N	Mean	SD	10%	25%	50%	75%	90%
Price ($)	Control	1200	49.02	63.62	12.15	21.16	34.66	57.20	92.17
	Treatment	1200	46.87	81.55	11.59	18.47	31.16	50.06	86.91
MarketCap ($ billion)	Control	1200	8.42	20.10	0.76	1.36	2.77	7.00	18.07
	Treatment	1200	6.47	34.86	0.60	0.80	1.42	3.56	9.11
DollarVolume ($ million)	Control	1200	69.48	132.87	4.14	9.86	25.81	71.78	168.24
	Treatment	1200	60.06	277.91	2.63	4.93	13.24	39.49	113.75
PSM Score	Control	1200	0.53	0.15	0.28	0.44	0.57	0.65	0.70
	Treatment	1200	0.54	0.16	0.28	0.44	0.60	0.68	0.71

where $metrc_{i,t}$ is the liquidity or trading variable of stock i on day t. α_i is the stock fixed effects. $NasdaqStock_{i,t}$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. $After_{i,t}$ is a dummy variable that equals one on and after October 24, 2016, and equals zero otherwise. Standard errors are clustered at the stock level.

The same identification strategy applies to the analysis of SIP glitch events. For all three SIP glitch events, it is either the Nasdaq-SIP or the NYSE-SIP that experienced a technical glitch, not both. So when NYSE-SIP experienced a glitch, we expect liquidity and trading in NYSE-listed stocks to be more affected than that in Nasdaq-listed stocks and vice versa when Nasdaq-SIP experienced a glitch. Importantly, as Figure A2, A3 and A4 in Appendix B shows, during all three SIP glitch events, direct feeds were largely unaffected and there is no market-wide trading halt. So the identified effect is not due to trade disruption as in the Nasdaq "Flash Freeze" event on August 22, 2013.

For each SIP glitch event, we use the same matching approach as above to construct a matched sample of 200 treated stocks (e.g., Nasdaq-listed stocks when Nasdaq-SIP experiences a glitch) and the same number of control stocks (e.g., NYSE-listed stocks when Nasdaq-SIP experiences a glitch). Table 5 reports the propensity matching score results and shows the matching is quite successful with a similar distribution of the two sample stocks across the matching variables and the final propensity score. The sample period for each event covers from 30 minutes before the start of the glitch until the end of the glitch. Pooling all three SIP-glitch events, we run the following standard DiD regression below:

$$metric_{i,d,t} = \alpha_{i,d} + \beta After_{i,d,t} + \gamma After_{d,t} \times Treated_{i,d} + \epsilon_{i,d,t}. \tag{5}$$

where $metrc_{i,d,t}$ is the liquidity or trading metric of stock i on event day d during the 30-second time interval t. $\alpha_{i,d}$ is the stock-day(event) fixed effect. $Treated_{i,d}$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock on January 3, 2013, and NYSE-listed stock on October 30, 2014, and August 12, 2019. $After_{d,t}$ is a dummy variable that equals one after the glitch starts, and equals zero otherwise. Standard errors are clustered at the stock-day(event) level. It is perhaps worth noting that the three SIP glitch events happened at different intraday periods and on both the Nasdaq-SIP and NYSE-SIP. Hence, pooling all three events in the regression helps alleviate some concerns about a possibly imperfect matching.

5.2 Nasdaq Venues vs. Other Venues

For the SIP upgrade analysis, our second identification strategy utilizes the unique geography of SIPs and exchanges in the U.S. equities market (See Figure 1). For a trader who subscribes to the SIP feeds, the total latency of receiving a message from one exchange consists of two parts: the traveling latency of the message, both from the exchange to SIP and then from SIP to the trader, and the processing latency of the message at the SIP. As an example, consider trading in Nasdaq-listed stocks. Messages from all exchanges have to be reported to and processed at Nasdaq-SIP. For a trader (labeled as "Trader A") who is located at Carteret, NJ where Nasdaq venues are located, the total latency of SIP will be dominated by the processing latency due to small traveling latency from Nasdaq venues to Nasdaq-SIP and from Nasdaq-SIP to the trader as they are all located at the same geographical location. In contrast, for a trader (labeled as "Trader B") who is located at Mahwah, NJ (where NYSE venues are located), a message from NYSE venues (at Mahwah, NJ) has to first travel to Nasdaq-SIP (at Carteret, NJ), be processed there and then travels back from Nasdaq-SIP (at Carteret, NJ) to the trader (at Mahwah, NJ).

After the Nasdaq-SIP speed upgrade which reduces the processing latency of the Nasdaq-SIP from 350 microseconds to less than 20 microseconds, the total SIP latency for Trader A is only about 20 microseconds. While for Trader B, the total latency will be 20 microseconds of processing latency plus a round trip traveling latency between Carteret, NJ and Mahwah NJ, which is roughly 560 microseconds through optical fiber and not affected by the reduction of the Nasdaq-SIP processing latency. In other words, the speed-up of Nasdaq-SIP will affect Trader A more than Trader B and thus trading on Nasdaq venues more than that on NYSE venues.[27]

To implement the identification strategy based on geographical latency, we estimate the following regression specification:

$$\begin{aligned} metric_{i,t,e} = \alpha_{i,e} &+ \beta After_t \\ &+ \gamma_1 After_t \times NasdaqStock_i \\ &+ \gamma_2 After_t \times NasdaqVenue_e \\ &+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e} \end{aligned} \tag{6}$$

where $metric_{i,t,e}$ is the liquidity or trading variable for stock i, traded on exchange e and on day t. $\alpha_{i,e}$ controls for the stock-venue fixed effects, which is potentially important as a stock can have a fixed trading pattern on a particular exchange. For example, although Nasdaq-listed stocks can trade at any of the 16 exchanges, opening and closing auctions are only held at Nasdaq, the listing exchange. Besides, for NYSE-listed stocks, the NYSE has DMMs (Designated Market Makers) who have mild obligations in maintaining liquidity. Last, exchanges have different fee schedules for stocks listed on its venue compared to stocks otherwise.[28] $After_{i,t}$ is dummy variable that

[27]For the sake of the example, here we assume that the two traders, Trader A and Trader B, are located at Carteret, NJ, and Mahwah, NJ, respectively so that the former experiences least geographical latency while the latter the most. However, in reality, traders' locations vary. If both traders are located at Secaucus, NJ, for example, the impact of the reduction of the Nasdaq-SIP processing latency will differ less for them due to the traveling latency of receiving the SIP feeds.

[28]See, e.g., https://www.nasdaqtrader.com/Trader.aspx?id=PriceListTrading2 for the fee schedule of Nasdaq

equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_i$, is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. $NasdaqVenue_i$ is a dummy variable that equals one if the variable is computed specifically on exchange e, and equals zero otherwise.[29] So the coefficient γ_3 captures the triple-difference-in-difference effect, that is, the change in the cross-venue difference for Nasdaq-listed stocks versus NYSE-listed stocks after the Nasdaq-SIP speed upgrade. Standard errors are clustered at the stock level. In Appendix A.4, as a robustness check, we use an alternative specification where we control for the stock-date fixed effects, instead of stock-venue fixed effects and the results are qualitatively the same.

6 Results

We next take the identification strategies developed in the previous section to the data and examine the role of consolidated feeds in the current U.S. equities market. We first focus on the event of a speed upgrade to the consolidated feed and examine its impact on market liquidity. Then we turn to events where the consolidated feed is disrupted due to a technical glitch.

6.1 Speed of the Consolidated Feed and Market Liquidity

6.1.1 NBBO and Nasdaq-BBO Dislocations

Before examining the impact of a faster consolidated feed on market liquidity, we first look at to what extent it makes its prices more reliable compared with direct feeds. To do so, we follow Bartlett and McCrary (2019) and construct two versions of NBBO, one based on the SIP Timestamp ("SIP-NBBO") and the other based on the Participant Timestamp[30] ("Participant-NBBO"). Dislocations between the two NBBOs occur when the SIP-NBBO lags and deviates from the pre-

venues.

[29]We only include four exchanges NYSE-Arca, Nasdaq, BZX and EDGX as they all adopt a maker-taker model. Moreover, the NYSE is excluded as it has only started to trade Nasdaq-listed stocks since April 2018.

[30]Participant Time in the SIP feeds represents the time when a quote update or trade is registered at the exchange's matching engine and thus suffers no processing latency at and traveling latency from/to the SIP.

vailing Participant-NBBO.[31] After the Nasdaq-SIP upgrade, which significantly reduces the processing latency of the SIP feed for Nasdaq-listed stocks, we expect them to experience fewer and shorter NBBO dislocations relative to NYSE-listed stocks.

In addition to market-wide NBBO, we look at the Nasdaq exchange in particular and construct its own BBO, again, based on SIP Timestamp ("SIP-Nasdaq-BBO") and Participant Timestamp ("Participant-Nasdaq-BBO") respectively. Recall that both the Nasdaq-SIP and the Nasdaq exchange are located at Carteret, NJ. So for Nasdaq-listed stocks, SIP feed reported from the Nasdaq exchange suffers barely any traveling latency but mainly processing latency at the Nasdaq-SIP. As a result, after the Nasdaq-SIP upgrade, which reduces its processing latency by over 90% to around 20 microseconds, we should see even fewer and shorter dislocations between the two Nasdaq-BBOs in Nasdaq-listed stocks relative to NYSE-listed stocks.

To visualize the impact of the Nasdaq-SIP upgrade on NBBO and Nasdaq-BBO dislocations, we compute, for each stock-day combination in our matched sample, two statistics: the number of and the median duration of dislocations for Nasdaq-listed versus NYSE-listed stocks. Figure 3 plots the cross-section average of the two statistics. As expected, after the upgrade, Nasdaq stocks saw a significant reduction in the median duration of both NBBO and Nasdaq-BBO dislocations. For example, the median duration of the Nasdaq-BBO dislocations drops from 200 microseconds to around only 10 microseconds, which implies that for traders who use SIP feeds for their trading in Nasdaq-listed stocks on the Nasdaq exchange, prices they receive become much more up-to-date. As for the number of NBBO and Nasdaq-BBO dislocations, we see a similar pattern: Both drop more for Nasdaq-listed stocks relative to NYSE-listed stocks after the upgrade. In addition, consistent with our expectation, the drop in Nasdaq-BBO dislocations is more pronounced. Quantitatively, the number of NBBO (Nasdaq-BBO) dislocations for Nasdaq-stocks drops by more than 1013 (1354) than NYSE-stocks, or roughly 18.72% (28.43%) relative to the mean level of Nasdaq-

[31] It should be noted that the dislocations we identify here do not necessarily represent true ones experienced by a consolidated feed subscriber compared with a direct feed subscriber. First, even for a direct feed subscriber, she will experience traveling latency and processing latency, although they can be smaller due to better networks (e.g., microwaves) and better hardware. Second, the actual latency a SIP subscriber experiences depends on its own location, which might be far away from SIPs, making the SIP prices she receives even more stale.

Figure 3. SIP Speed Upgrade and NBBO Dislocations. Following Bartlett and McCrary (2019), we construct two versions of NBBOs, one based on the SIP Timestamps (SIP-NBBO) and the other based on the Participant Timestamp (Participant-NBBO), and then identify the dislocations between the two NBBOs. By the same token, we construct two BBOs on the Nasdaq exchange, one based again on the SIP Timestamp and the other on the Participant Timestamp.

(a) This figure plots the median duration and number of NBBO dislocations for NYSE-listed Stocks and Nasdaq-listed stocks respectively. The time series plotted is the cross-section average of all sample stocks.

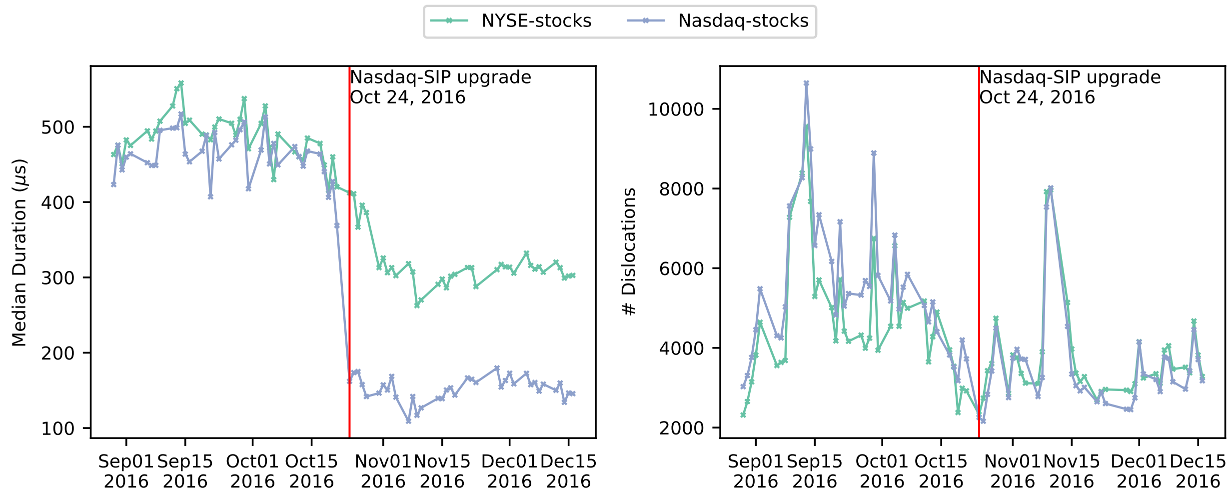

(b) This figure plots the median duration and the number of Nasdaq-BBO dislocation dislocations for NYSE-listed Stocks and Nasdaq-listed stocks respectively. The time series plotted is the cross-section average of all sample stocks.

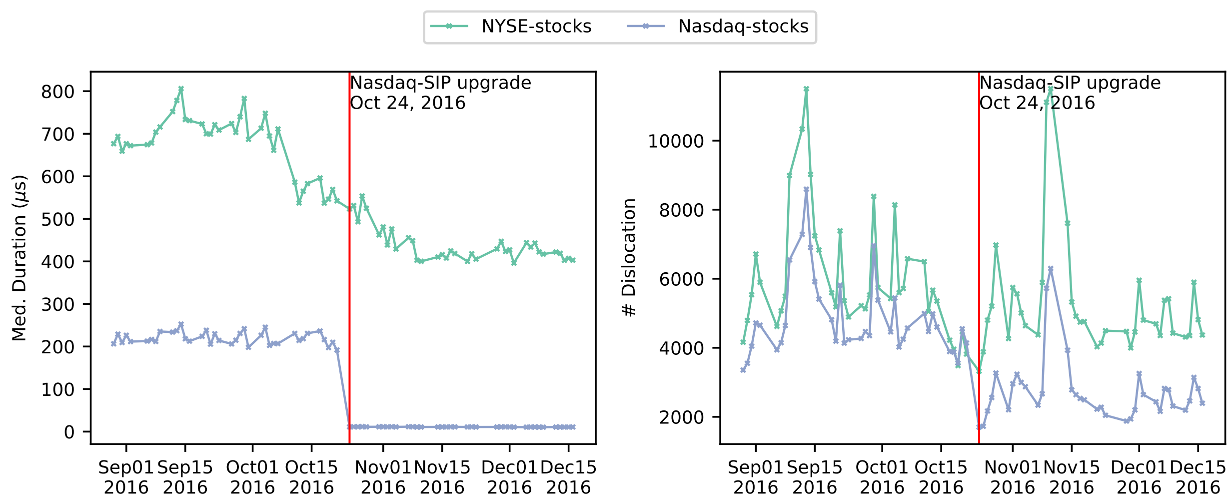

stocks before the upgrade.

Summing up, the foregoing results of the NBBO and Nasdaq-BBO dislocations show that after the Nasdaq-SIP upgrade, pricing information contained in the consolidated feed for Nasdaq-listed stocks, especially on the Nasdaq exchange, experiences much less latency and becomes more reliable relative to NYSE-listed stocks. Whether such an improvement in the consolidated feed prices can translate to better market liquidity or have a significant impact on market participants' trading behavior remains an empirical question that we will tackle below.

6.1.2 Market-Wide Liquidity

How would a faster consolidated feed affect market liquidity? We examine the impact of the Nasdaq-SIP upgrade on market-wide liquidity by estimating the DiD regression specified in Equation 4. As detailed in Section 5.1, the upgrade only reduces the latency of the consolidated feed for Nasdaq-listed stocks, not NYSE-listed stocks. So if there is any effect, we should see our liquidity measures change significantly for Nasdaq-listed stocks relative to NYSE-listed ones.

We report the estimation results for liquidity variables in Table 6. It shows that the coefficient on the interaction term *After* × *NasdaqStock*, which captures the DiD effect, is statistically significant for relative quote spread (*RQS*) and relative effective spread (*RES*). Specifically, these two spread measures increase by 0.54 and 0.15 basis points respectively for Nasdaq-listed stocks relative to NYSE-listed stocks. In terms of economic magnitudes, the increases are about 6.32% (≈0.54 ÷ 8.55) and 4.76% (≈0.15 ÷ 3.15) compared to their unconditional means across all sample stocks.

The relative realized spread (*RRS*) does not have a statistically significant change. The increase in the relative effective spread is almost completely offset by a corresponding increase in relative price impact (*RPI*). Specifically, the relative price impact increases by 0.13 or 5.53% (≈0.13 ÷ 2.35) relatively after the SIP upgrade. For the rest of our liquidity measures, the results are insignificant: We do not see a significant change in the NBBO depth (*Depth*) or market volume (*Vlm*). A faster consolidated feed seems to only affect the price of market liquidity, not the quantity of it.

What might cause the worsening of market liquidity? One possible explanation is that as the speed advantage of direct feeds over consolidated feeds greatly vanishes as the latter becomes faster, HFTs now face more competition from slower traders (e.g., buy-side execution algorithms). As a result, HFT market makers have to widen their spreads to compensate for higher adverse selection costs as evidenced by larger price impacts following the upgrade.

Table 6. Market-Wide Liquidity Impact of a Faster Consolidated Feed. This table shows the estimation results from the DiD regression specified below:

$$metric_{i,t} = \alpha_i + \beta After_t + \gamma After_t \times NasdaqStock_i + \epsilon_{i,t}$$

where $metrc_{i,t}$ is the liquidity or trading variable of stock i on day t. α_i is the stock fixed effects. $After_{i,t}$ is a dummy variable that equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_{i,t}$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact, and relative realized spread respectively, and are all in basis point. *Depth* is the NBBO depth in thousand dollars. *Vlm* is trading volume in million dollars.

	RQS	RES	RPI	RRS	Depth	Vlm
After	0.63^{***}	0.11^{***}	0.14^{***}	−0.03	-8.97^{**}	25.93^{***}
	(0.07)	(0.03)	(0.02)	(0.03)	(3.72)	(2.94)
After x NasdaqStock	0.54^{***}	0.15^{**}	0.13^{***}	0.02	−7.51	−4.02
	(0.20)	(0.06)	(0.04)	(0.05)	(13.64)	(6.25)
R^2 (%)	2.13	0.27	0.65	0.00	0.21	0.95
N	42900	42900	42900	42900	42900	42900
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

The above explanation is supported by the impact of a faster consolidated feed on trading. In Table 7, we report the results of the DiD estimation with our trading variables. There are several findings worth discussing. First, the coefficient on the interaction term $After \times NasdaqStock$ is significantly positive for our two high-frequency market-making proxies, *Cancel/Trade* and *Order/Trade*. The economic magnitudes are relatively large: The DiD increases in the two proxies are 3.09 and 6.13 respectively, about 12.99% (≈3.09 ÷ 23.79) and 16.11% (≈6.13 ÷ 38.04) relative to their unconditional means. The estimation results suggest that, following a faster consolidated feed, there is a significant increase in overall quoting activity for Nasdaq-listed stocks relative to NYSE-listed stocks.

Second, we find that the coefficient on the interaction term $After \times NasdaqStock$ is significantly positive for the share of ISO trades (*ISOShr*), ISO orders are used to bypass the Reg-NMS order protection rule to trade behind-the-top depth on a target exchange by sweeping through top-of-book depth across all other exchanges with better prices. Thus, as argued by Chakravarty, Jain, Upson, and Wood (2012), they are mainly used by informed institutional traders to increase execution

Table 7. Market-Wide Trading Impact of a Faster Consolidated Feed. This table shows the estimation results from the DiD regression specified below:

$$metric_{i,t} = \alpha_i + \beta After_t + \gamma After_t \times NasdaqStock_i + \epsilon_{i,t}$$

where $metrc_{i,t}$ is the liquidity or trading variable of stock i on day t. α_i is the stock fixed effects. $After_{i,t}$ is a dummy variable that equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_{i,t}$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016. *ISOShr* is trade volume via inter-market sweep order (ISO) as a fraction of total trade volume. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	-2.16^{***}	0.09	-8.32^{***}	-14.21^{***}
	(0.16)	(0.10)	(0.44)	(0.98)
After x NasdaqStock	1.53^{***}	0.18	3.09^{***}	6.13^{***}
	(0.24)	(0.14)	(0.53)	(1.08)
R^2 (%)	1.80	0.12	16.16	12.88
N	42900	42900	42900	42900
Stock F.E.	Yes	Yes	Yes	Yes

speed and capture larger depth. A faster consolidated feed with more up-to-date price information might encourage buy-side execution algorithms that use them as market data input to submit more ISO orders. However, we caution that the economic magnitude of the increase in ISO volume share is rather small and about 4.37% (≈1.53 ÷ 35.03) to its unconditional mean.

The coefficient on odd-lot share is not statistically significant, though the empirical evidence on whether odd-lot trades are informed or not is mixed. While O'Hara, Yao, and Ye (2014) show that odd-lot trades are more informative than mixed- and round-lot trades, Upson and Johnson (2017) make a further distinction between the odd-lot trades resulting from odd-lot resting limit orders and those directly from odd-lot marketable orders and show the latter is not more informed than mixed-lot and round-lot marketable orders.

Table 8. Exchange-Specific Liquidity Impact of a Faster Consolidated Feed. This table shows results from the triple-difference-in-difference regression

$$metric_{i,t,e} = \alpha_{i,e} + \beta After_t + \gamma_1 After_t \times NasdaqStock_i + \gamma_2 After_t \times NasdaqVenue_e$$
$$+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e}$$

where $metric_{i,t,e}$ is the liquidity or trading metric for stock i, traded on exchange e and on day t. $\alpha_{i,e}$ controls for the stock-venue fixed effects. $After_{i,t}$ is dummy variable that equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_i$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. $NasdaqVenue_i$ is dummy variable that equals one if the metric is computed based on exchange e, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016. All measures have been defined above in Table 6. But note that now they are measures specific to an exchange. For example, *Depth* of an exchange will be its time-weighted dollar depth at the NBBO.

	RQS	RES	RPI	RRS	Vlm	Depth
After	3.87***	0.15***	0.21***	−0.06***	1.71***	−1.05**
	(0.54)	(0.02)	(0.03)	(0.02)	(0.20)	(0.44)
After x NasdaqStock	0.55	0.09*	0.12*	−0.03	−0.02	−1.16
	(0.88)	(0.05)	(0.07)	(0.04)	(0.49)	(2.01)
After x NasdaqVenue	−3.35***	−0.03***	0.01	−0.03**	2.73***	1.44***
	(0.54)	(0.01)	(0.02)	(0.01)	(0.22)	(0.27)
After x NasdaqStock x NasdaqVenue	−0.25	0.06***	0.09***	−0.03	1.78**	0.56
	(0.80)	(0.01)	(0.03)	(0.03)	(0.74)	(1.27)
R^2 (%)	0.94	1.05	0.59	0.07	1.02	0.08
N	171600	171597	171597	171597	171597	171600
Stock-Venue F.E.	Yes	Yes	Yes	Yes	Yes	Yes

6.1.3 Exchange-Specific Liquidity

In the previous section, we examine the impact of a faster consolidated feed on market-wide liquidity. Would the impact be different across different exchanges? As we illustrate in Section 5.2, the consolidated feed in the US equities has a unique geography. Specifically, both Nasdaq exchanges and the Nasdaq-SIP are located at Carteret, NJ. The geographical proximity of the two means that messages from Nasdaq exchanges experience the smallest traveling latency to the Nasdaq-SIP and thus processing latency at the SIP makes up the largest component of total latency. Consequently, when the Nasdaq-SIP upgrade substantially cuts down its processing latency, messages from Nasdaq exchanges are likely to see the most significant reduction in overall latency. To explore it, we employ the triple-DiD regression specified in Equation 6, using exchange-specific metrics as dependent variables.

Triple-DiD estimation results for liquidity variables are reported in Table 8, and they show that the coefficient on the interaction term *After* × *NasdaqStock* × *NasdaqVenue*, which captures the triple-DiD effect, is positively significant for *RES* and *RPI*. In the DiD results above, we document that both *RES* and *RPI* increase for Nasdaq-listed stocks relative to NYSE-listed stocks after the Nasdaq-SIP upgrade. The triple-DiD results further suggest that the effect is larger on the Nasdaq exchange relative to other exchanges, supporting our hypothesis that a faster consolidated feed leads to larger price impacts or more informed trading. As detailed above, while there is an overall reduction of the consolidated feed's latency for Nasdaq-listed stocks, the relative magnitude is much more significant on Nasdaq exchanges. Thus prices from the consolidated feed for Nasdaq-listed stocks traded on Nasdaq exchanges become much more up-to-date after the upgrade, encouraging the trading of slower traders like buy-side execution algorithms using the consolidated feed.

Triple-DiD estimation results for trading variables reported in Table 9 show the same increase in the high-frequency quoting activities. The coefficient on the triple interaction term for the two high-frequency quoting proxies, *Cancel/Trade* and *Order/Trade*, is significantly positive and has a similar magnitude as in the DiD results above. In contrast, the coefficient on the interaction term *After* × *NasdaqStock* is insignificant, suggesting that the increased high-frequency quoting activity is only observed in Nasdaq-listed stocks on the Nasdaq exchange.

In addition to the "algorithmic trading" channel where a faster consolidated feed results in more competition between HFT market makers and informed algorithmic traders and thus more high-frequency quoting activities, an alternative explanation is the "quote stuffing" channel. A larger SIP capacity forces HFTs who implement a quote stuffing strategy to send more "garbage" messages and thus increase the overall message traffic. The triple-DiD results favor the first, benign channel. If HFTs' goal is to slow down the Nasdaq-SIP, it shouldn't matter to which exchange they send "garbage" messages. As a matter of fact, to avoid violating exchange messaging policy[32], it

[32]For example, Nasdaq has an Excessive Messaging Policy that discourages excessive order activity away from the

Table 9. Exchange-Specific Trading Impact of a Faster Consolidated Feed. This table shows results from the triple-difference-in-difference regression

$$metric_{i,t,e} = \alpha_{i,e} + \beta After_t + \gamma_1 After_t \times NasdaqStock_i + \gamma_2 After_t \times NasdaqVenue_e + \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e}$$

where $metric_{i,t,e}$ is the liquidity or trading metric for stock i, traded on exchange e and on day t. $\alpha_{i,e}$ controls for the stock-venue fixed effects. $After_{i,t}$ is dummy variable that equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_i$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. $NasdaqVenue_i$ is dummy variable that equals one if the metric is computed based on exchange e, and equals zero otherwise. Standard errors are clustered at the stock level. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. Note that we exclude stocks involved in SEC's Tick Size Pilot Program. The sample period is from August 29 to December 16, 2016. All measures have been defined above in Table 7. But note that now they are measures specific to an exchange. For example, *Cancel/Trade* of an exchange will be the cancel-to-trade ratio on its own venue.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	-1.95^{***}	0.28^{**}	-7.69^{***}	-10.86^{***}
	(0.22)	(0.12)	(0.54)	(0.75)
After x NasdaqStock	1.18^{***}	0.21	0.54	0.07
	(0.31)	(0.16)	(0.78)	(1.16)
After x NasdaqVenue	-0.39^{***}	0.42^{***}	-1.08^{***}	-1.47^{**}
	(0.12)	(0.07)	(0.37)	(0.58)
After x NasdaqStock x NasdaqVenue	0.57^{***}	-0.35^{***}	4.33^{***}	5.85^{***}
	(0.19)	(0.11)	(0.59)	(0.99)
R^2 (%)	1.01	0.29	2.26	0.19
N	171597	171597	171595	171595
Stock-Venue F.E.	Yes	Yes	Yes	Yes

is a better strategy to spread "garbage" messages across different exchanges. So if we see the two high-frequency quoting proxies increase more on the Nasdaq exchange than others, it is more likely to be caused by more normal market-making activities.

The coefficient on the triple-interaction term is significantly positive for *ISOShr* while negative for *OddlotShr*. Recall that ISO orders are used by traders to trade behind-the-top depth at a target exchange. So if a faster consolidated feed makes the Nasdaq exchange more likely to be the target exchange of ISO traders, then the share of ISO trades on the venue might increase. However, it is worth mentioning that we are cautious about the triple-DiD results for ISO volume shares. In Section A.3 of the appendix, we shorten the sample period to two months and the coefficients turn

NBBO. Specifically, member firms that exceed a "Weighted Order-to-Trade Ratio" of 100:1 pay a fee on the orders that cause the firm to exceed the threshold.

statistically insignificant.

In summary, both the DiD and triple-DiD results show that a faster consolidated feed leads to slightly worse market liquidity as evidenced by higher spreads and higher price impacts. In addition, they lend support to the explanation that as slow traders such as buy-side execution algorithms become faster, HFT market makers face more competition. As a result, they engage in more frequent quote updating activities and charge larger spreads to offset higher adverse selection costs.

6.1.4 Direct Evidence of More Efficient Buy-Side Executions

In the previous section, we show that market liquidity worsens when the consolidated feed becomes faster and it might be due to an increase in informed algorithmic trading from the buy-side. Below, we provide two direct evidence supporting our conjecture.

Number of sub-optimally routed trades The first measure we use is the number of sub-optimally routed trades. In the U.S. equities market, exchanges adopt different fee models. Broadly, there are two types: "Maker-Taker" exchanges (e.g., Nasdaq, NYSE, NYSE Arca, BZX, and EDGX) provide rebates to makers (i.e., liquidity providers) and charge fees from takers (i.e, liquidity consumers); "Taker-Maker" exchanges or inverted exchanges (e.g., BX, BYX, and EDGA) implements the opposite pricing model. They provide rebates to takers and charge fees from makers. If a trader routes her orders to a "Maker-Taker" exchange, that is, paying a fee to take liquidity, while there's liquidity at the same price on an inverted exchange where they get paid to take liquidity, that's a clear sign of sub-optimal routing and trading inefficiency.

To pinpoint these trades, we construct the best bids and offers for both exchange types using the Participant Timestamp and SIP Timestamp, respectively. We then identify trades executed on Maker-Taker exchanges at prices not reflected on the inverted exchange when using the SIP Timestamp, but available when referencing the Participant Timestamp. It suggests that traders, using the slower consolidated feed, are making decisions that result in sub-optimal trade routing.

Figure 4. Number of Sub-Optimally Routed Trades. This figure plots the daily number of sub-optimally routed trades for NYSE-listed stocks and Nasdaq-listed stocks in our sample respectively. A sub-optimally routed trade refers to a trade executed on a Maker-Take exchange when the same price is available on an inverted exchange based on the Participant Timestamp but not the SIP Timestamp. The vertical line represents the speed upgrade to the Nasdaq-SIP on October 24, 2016. The table below breaks down the average daily number of such trades for the two stock types and two periods respectively.

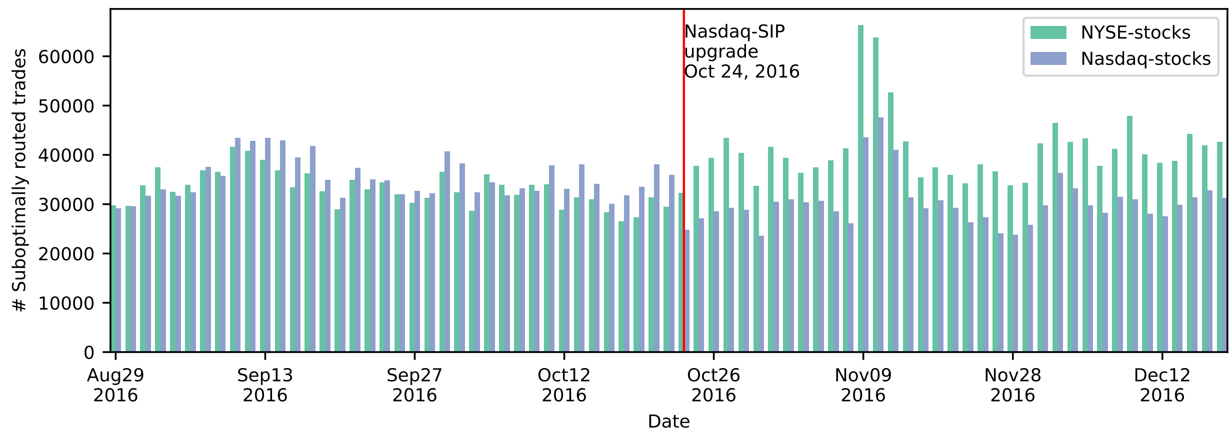

	NYSE-stocks	Nasdaq-stocks	Difference
Before Upgrade	33052	35266	2214
After Upgrade	41087	30241	-10845
Difference	8034	-5025	-13059

Figure 4 plots the number of sub-optimally routed trades for NYSE-listed stocks and Nasdaq-listed stocks respectively. It shows that there's a noticeable shift in the number of sub-optimally routed trades after the Nasdaq-SIP upgrade: NYSE-listed stocks see an increase, whereas Nasdaq-listed stocks witness a decrease. The table below the graph provides a numerical breakdown. Specifically, before the upgrade, there were 2,214 fewer sub-optimal trades for NYSE-listed stocks than for Nasdaq-listed stocks. However, after the upgrade, the difference flipped sign and changed to -10,845 trades. It means that post-upgrade, NYSE-listed stocks saw more sub-optimal trades compared to Nasdaq-listed stocks. The difference-in-difference result is -13,059, which underscores the contrasting impact of the upgrade on the two stock categories. The results clearly suggest that as the consolidated feed becomes faster, buy-side execution algorithms make fewer sub-optimally routed trades.

Fill time of inverted liquidity Another measurement we analyze is the duration it takes to fulfill an order on an inverted exchange (i.e., the fill time). The idea is straightforward: if the SIP upgrades prove beneficial, traders utilizing the consolidated feed should identify available liquidity more promptly after the upgrade, which should in turn decrease the fill time.

As we don't have order-level limit order book data tracking the life cycle of all submitted orders, we instead use the top-of-book updates on the inverted exchanges. Specifically, we identify cases where the best bid and ask on an inverted exchange improve, that is, the best bid increases or the best ask decreases and it is immediately followed by a trade execution.

We decide to only keep trades with a fill time that is larger than 20 microseconds and smaller than 10 milliseconds. These trades should capture market participants at various locations of the latency spectrum. For HFTs using co-location and having access to direct feeds, they should be on the lower end of the spectrum, for buy-side execution algorithms, their responses should be near the other end of the spectrum. In Aquilina, Budish, and O'Neill (2021), the authors analyze the U.K. equities market based on a sample in 2015 and use 29 microseconds as the lower bound for the response time of the fastest traders. We use a threshold similar to their choice.

Figure 5 illustrates cumulative histograms for the fill time of BBO-improved inverted liquidity, for NYSE stocks and Nasdaq stocks, before and after the Nasdaq-SIP upgrade. It shows that the two cumulative histograms, one for NYSE-listed stocks and another for Nasdaq-listed stocks, follow a very similar trajectory before the upgrade, implying comparable fill times between the two stock types. After the upgrade, there's a slight divergence in the density of the two, especially in the middle fill time ranges around 1 millisecond.

The table below the cumulative histogram graphs breaks down the fill times for NYSE-listed and Nasdaq-listed stocks into specific quantiles (from Q1 to Q90) for both before and after the upgrade. The "Difference" rows capture the difference in fill times between NYSE and Nasdaq-listed stocks for each quantile. From the table, one can observe that the differences between NYSE and Nasdaq-listed stocks in terms of fill times are minimal in the earlier quantiles but start to expand as we move to the higher quantiles or to the range around 500 microseconds, especially

Figure 5. Fill Time of Liquidity on Inverted Exchanges. This figure plots the cumulative histogram of the fill time of BBO-improving liquidity on inverted exchanges, for NYSE-listed and Nasdaq-listed stocks in our sample. The table breaks down the fill times for NYSE-listed and Nasdaq-listed stocks into specific latency quantiles.

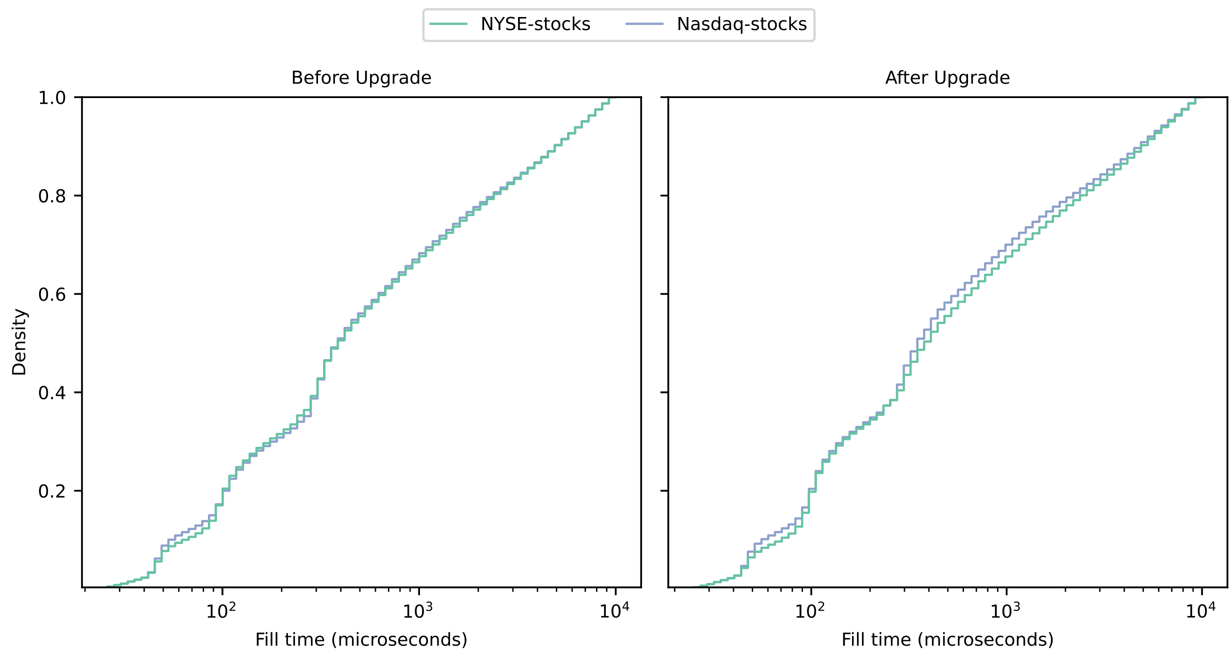

	Quantile	Q1	Q5	Q10	Q20	Q30	Q40	Q50	Q60	Q70	Q80	Q90
Before upgrade	NYSE-stocks	32	49	67	107	179	309	408	682	1265	2528	5221
	Nasdaq-stocks	32	48	57	108	190	311	402	666	1212	2454	5200
	Difference	0	-1	-10	1	11	1	-6	-16	-53	-74	-20
After upgrade	NYSE-stocks	31	49	73	106	153	294	403	670	1252	2557	5200
	Nasdaq-stocks	31	47	59	105	148	289	366	578	1067	2264	4989
	Difference	0	-1	-13	-1	-4	-5	-36	-91	-185	-292	-211

after the upgrade. This suggests that while the SIP upgrade might not have significantly affected the majority of trades, it does have an impact on the tail-end, or the trades with longer fill times. The results are consistent with our conjecture. A faster consolidated feed should not affect low-latency traders such as HFTs because they use direct feeds instead. But for buy-side execution algorithms, a faster consolidated feed can potentially benefit them.

In summary, both analyses collectively present evidence that a faster consolidated feed, brought by the Nasdaq-SIP upgrade, provides benefits to slow traders such as buy-side execution algorithms, optimizing their trade execution and efficiency. Faced with higher competition and larger price impacts, HFT market makers respond by widening their spreads, deteriorating market liquid-

Table 10. Market Liquidity Impact of a Disrupted Consolidated Feed. This table shows regression results from the following difference-in-difference regression:

$$metric_{i,d,t} = \alpha_{i,d} + \beta After_{i,d,t} + \gamma After_{d,t} \times Treated_{i,d} + \epsilon_{i,d,t}.$$

where $metrc_{i,d,t}$ is the liquidity or trading metric of stock i on event day d during the 30-second time interval t. $\alpha_{i,d}$ is the stock-day fixed effect. $Treated_{i,d}$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock on January 3, 2013, and NYSE-listed stock on October 30, 2014, and August 12, 2019. $After_{d,t}$ is a dummy variable that equals one after the glitch starts, and equals zero otherwise. Standard errors are clustered at the event-stock level. The sample stocks consist of 1200 treated stocks (i.e., affected by the SIP glitch) matched with the same number of control stocks on price, market capitalization, trading volume, and industry. The sample period is between 30 minutes before the start of the SIP glitch and the end of it. *RQS*, *RES* and *RRS* stands for relative quoted spread, effective spread, and realized spread respectively, and are all in basis point. *Vlm* is dollar volume in thousand dollars. *DepthNBBO* is NBBO depth in thousand dollars. *Depth5Lvl* is cumulative depth across the five best price levels of the order book.

	RQS	RES	RRS	Vlm	DepthNBBO	Depth5Lvl
After	0.35^{*}	−0.04	−0.02	2.85	−3.16	-35.01^{***}
	(0.20)	(0.05)	(0.08)	(1.84)	(2.13)	(4.77)
After x Treated	0.59^{**}	0.37^{***}	0.96^{***}	-5.93^{***}	-17.91^{***}	-64.85^{***}
	(0.25)	(0.07)	(0.10)	(2.14)	(6.55)	(23.42)
R^2 (%)	0.22	0.14	0.31	0.01	0.09	0.77
N	235564	129309	129309	235564	235564	235564
Stock-Event F.E.	Yes	Yes	Yes	Yes	Yes	Yes

ity. As a consequence, small, uninformed traders like retail might be harmed.

6.2 Consolidated Feed Disruption and Market Liquidity

Following examining the market impact of a faster consolidated feed, now we turn to SIP glitches and investigate what happens to the market when the consolidated feed is disrupted.

To study the impact of SIP glitches on market liquidity, we run the DiD regression specified in Equation 5 where we pool all three major SIP glitch events detailed in Table 1. The identification strategy is identical to what we use in the analysis of SIP upgrade events. Recall that during all three SIP glitch events, only one of the two SIPs is affected. Thus stocks whose consolidated feed is (not) affected serve naturally as the treatment (control) group. So if SIP glitches have a negative impact on market liquidity, we should see our liquidity measures worsen more for treated stocks than control stocks.

Table 10 reports the regression results. It shows that stocks with their consolidated feed disrupted due to a SIP glitch see their market liquidity worsen by all measures. Specifically, the coefficient on the interaction dummy *After* × *Treated*, which captures the DiD effect, shows that the relative quoted spread and relative effective spread of treatment stocks increase by 0.59 and 0.37 basis points more than control stocks respectively. In terms of economic magnitudes, they correspond to about 4.28% (≈0.59 ÷ 13.80) and 11.14% (≈0.37 ÷ 3.32) relative to their unconditional means.

In addition, a disrupted consolidated feed has a much more significant and negative impact on trading volume and order-book depth. Specifically, the trading volume of treatment stocks falls by nearly 17.81% (≈-5.93 ÷ 33.30) more than control stocks. As for the two order-book depth measures, they follow a similar pattern: NBBO depth and cumulative depth across the five best price levels drop by about 21.20% (≈-17.91 ÷ 84.50) and 13.80% (≈-64.85 ÷ 470.26) more for treatment stocks than control stocks respectively.

Note that the coefficient on the time dummy *After* is significantly positive for relative quoted spread and negative for cumulative depth across five best prices as well, indicating that market liquidity of control stocks is affected as well. The result is perhaps not surprising as trading in the two matched samples of stocks is correlated, either due to them being in the same market index or in the same industry. So when treatment stocks become illiquid, the illiquidity can quickly spread to the control stocks. Such illiquidity contagion through informationally correlated assets is well modeled in Cespa and Foucault (2014). An alternative explanation is that some HFTs or execution algorithms might cease their market-making or other trading activities for all stocks when faced with a market-wide data anomaly.

In addition to pooling all three SIP glitch events as above, we zoom in on the Nasdaq-SIP glitch event on January 3, 2013, and exploit its unique feature for further identification. The results show a similar pattern. The liquidity of stocks in dissemination channels that experienced a glitch earlier deteriorates first, followed by other stocks affected later. We provide detailed results in Appendix C.

7 Conclusion

In this paper, we examine the role of the consolidated feed in a fast and fragmented market. We zoom in to the U.S. equities market for an empirical study by examining exogenous shocks when its consolidated feed becomes faster or disrupted. The unique structure of the consolidated feed, one stream for Nasdaq-listed stocks and the other stream for NYSE-listed stocks, allows us to implement a clean difference-in-difference analysis based on a matched sample of Nasdaq-listed stocks and NYSE-listed stocks. The results show that a faster consolidated feed has a mild and adverse effect on market liquidity, as indicated by larger spreads. Higher price impacts and quoting activities suggest that the worsening of market liquidity might result from increased competition between high-frequency market makers and informed slow traders like buy-side execution algorithms as the latter becomes faster. Such an explanation is further supported by evidence of slow traders experiencing fewer sub-optimally routed trades and a shorter fill time of inverted liquidity. In addition, we document that when the consolidated feed becomes disrupted due to technical glitches, market liquidity, especially the market volume and order-book depth, worsens significantly.

Our results show that the consolidated feed matters and is a crucial component of today's market data infrastructure. While our study concentrates mainly on the aspect of speed – a significant factor – it's clear that more research is needed to fully evaluate the implications of potential regulatory changes to the consolidated feed including the integration of odd-lot quotes and the introduction of competing consolidators, among others.

Appendix

A Robustness Checks

Here we perform several robustness checks for our difference-in-difference estimation in Section 6.1.

A.1 Visual Evidence for Parallel Trend Assumption

Figure A1 plots the market-wide liquidity and trading metrics used in the DiD regression and presents visual evidence for the parallel trend assumptions. It shows that all variables for the treatment group (i.e., Nasdaq-listed stocks) evolved in a similar fashion to the control group (i.e., NYSE-listed stocks) before the event date, suggesting that the parallel trend assumption is supported.

A.2 Placebo Tests

We perform two placebo tests around the Nasdaq-SIP upgrade event as a robustness check to assess the validity of the parallel trends assumption and to rule out the possibility that observed treatment effects are due to chance or confounding factors. First, we shift the event date to September 26, 2016, and perform the DiD estimation for a two-month window before the actual event date. Second, we shift the event date to November 21, 2016, and perform the DiD estimation for a two-month window after the actual event date.

Pre-upgrade placebo tests Our placebo tests further substantiate the robustness of our baseline findings. Conducting difference-in-difference estimates with a shifted event date prior to the actual Nasdaq-SIP upgrade, all estimates, with the exception of *ISOShr*, turn out to be statistically insignificant, as evidenced in Table A1 and Table A2. The results strengthen the validity of our main

Table A1. Pre-Upgrade Placebo Test: Liquidity. This table presents DiD estimates using a shifted event date of September 26, 2016, instead of the actual Nasdaq-SIP upgrade on October 24, 2016. The focus is on liquidity metrics over a two-month window preceding the real event. Estimation details align with baseline specification in Table 6.

	RQS	RES	RPI	RRS	Depth	Vlm
After	0.15**	−0.01	−0.06**	0.06	−7.56**	0.51
	(0.06)	(0.05)	(0.03)	(0.04)	(3.66)	(2.13)
After x NasdaqStock	0.00	0.00	0.01	−0.01	6.65	−2.13
	(0.16)	(0.08)	(0.04)	(0.07)	(5.29)	(6.07)
R^2 (%)	0.10	0.00	0.05	0.02	0.12	0.00
N	21964	21964	21964	21964	21964	21964
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Table A2. Pre-Upgrade Placebo Test: Trading. This table presents DiD estimates using a shifted event date of September 26, 2016, instead of the actual Nasdaq-SIP upgrade on October 24, 2016. The focus is on trading metrics over a two-month window preceding the real event. Estimation details align with baseline specification in Table 7.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	0.99***	−0.07	−1.72***	−3.35***
	(0.18)	(0.09)	(0.30)	(0.59)
After x NasdaqStock	−0.84***	0.14	−0.24	0.69
	(0.27)	(0.14)	(0.39)	(0.72)
R^2 (%)	0.39	0.02	1.17	0.88
N	21964	21964	21964	21964
Stock F.E.	Yes	Yes	Yes	Yes

findings, confirming that the worsening of market liquidity is indeed attributable to the upgrade and not to pre-existing trends.

Post-upgrade placebo tests The robustness of our primary findings is further supported by post-upgrade placebo tests as reported in Table A3 and Table A4. Shifting the event date to a time after the actual Nasdaq-SIP upgrade, we find that all estimates for liquidity measures are statistically insignificant. As for the trading variables, while *Cancel/Trade* is statistically significant, its magnitudes are considerably smaller compared to the baseline results. This attenuation in magnitude may be due to a delayed reaction from market participants.

Table A3. Post-Upgrade Placebo Test: Liquidity. This table presents DiD estimates using a shifted event date of November 21, 2016, instead of the actual Nasdaq-SIP upgrade on October 24, 2016. The focus is on liquidity metrics over a two-month window after the real event. Estimation details align with baseline specification in Table 6.

	RQS	RES	RPI	RRS	Depth	Vlm
After	-0.27^{***}	-0.13^{***}	-0.22^{***}	0.09^{**}	2.23	-12.05^{***}
	(0.08)	(0.04)	(0.03)	(0.04)	(4.46)	(3.18)
After x NasdaqStock	0.27	0.08	0.06	0.02	−19.75	−4.55
	(0.22)	(0.07)	(0.05)	(0.06)	(15.77)	(6.40)
R^2 (%)	0.08	0.06	0.58	0.07	0.18	0.37
N	21201	21201	21201	21201	21201	21201
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Table A4. Post-Upgrade Placebo Test: Trading. This table presents DiD estimates using a shifted event date of November, 2016, instead of the actual Nasdaq-SIP upgrade on October 24, 2016. The focus is on trading metrics over a two-month window after the real event. Estimation details align with baseline specification in Table 7.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	-1.92^{***}	1.43^{***}	-2.24^{***}	-2.06^{***}
	(0.19)	(0.10)	(0.24)	(0.46)
After x NasdaqStock	0.43	-0.45^{***}	1.18^{***}	0.73
	(0.26)	(0.15)	(0.29)	(0.55)
R^2 (%)	2.30	4.73	2.36	0.70
N	21201	21201	21201	21201
Stock F.E.	Yes	Yes	Yes	Yes

A.3 Changing the Length of Event Window

In our baseline regression, we focus on the Nasdaq-SIP speed upgrade on October 24, 2016, and use a four-month window, two months before the event date and two months after. As a robustness check, here we shorten the window length to two months, one month before the event date and one month after, to mitigate the concern that we might capture confounding events other than the Nasdaq-SIP speed upgrade.

Table A5 and Table A6 report the estimation results of the DiD regression specified in Equation 4 for liquidity and trading respectively. While most results stay qualitatively the same as the baseline, we would like to mention two noticeable changes. On liquidity, the coefficient on the interaction term turns insignificant for *RES*, indicating that while the relative quoted spread in-

Table A5. DiD Regression with a Two-Month Event Window: Liquidity. This table shows the DiD estimation results with a two-month window around the event day for liquidity variables. Estimation details align with baseline specification in Table 6.

	RQS	RES	RPI	RRS	Depth	Vlm
After	0.69***	0.18***	0.28***	−0.10***	−6.32***	31.39***
	(0.07)	(0.05)	(0.03)	(0.04)	(2.25)	(3.69)
After x NasdaqStock	0.41**	0.11	0.09*	0.02	−1.49	−0.24
	(0.20)	(0.08)	(0.05)	(0.06)	(7.70)	(7.74)
R^2 (%)	2.12	0.32	1.35	0.06	0.11	1.45
N	21888	21888	21888	21888	21888	21888
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Table A6. DiD Regression with a Two-Month Event Window: Liquidity. This table shows the DiD estimation results with a two-month window around the event day for trading variables. Estimation details align with baseline specification in Table 7.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade	#Run/Vlm
After	−1.69***	−0.56***	−6.37***	−11.54***	−1.72***
	(0.18)	(0.10)	(0.42)	(0.94)	(0.36)
After x NasdaqStock	1.74***	0.32**	2.62***	5.42***	0.04
	(0.25)	(0.14)	(0.50)	(1.04)	(0.55)
R^2 (%)	1.07	0.60	10.64	9.45	0.83
N	21888	21888	21888	21888	21888
Stock F.E.	Yes	Yes	Yes	Yes	Yes

creases, traders adapt to finding better prices than NBBO. For example, traders can move to dark pools or access on-exchange hidden liquidity. Turning to trading variables, only the coefficient for *OddlotShr* changes and turns significant. However, its economic magnitude of 0.32% is rather small.

Table A7 and Table A8 report the estimation results of the triple DiD regression specified in Equation 6 for liquidity the trading metrics respectively. On liquidity, the coefficient for all variables is qualitatively the same as the baseline results. On trading, we have already noted in the main text that the coefficient on *ISOShr* and *OddlotShr* turn insignificant. Other key results for the two AT proxies stay qualitatively the same.

Table A7. Triple DiD Regression with a Two-Month Event Window: Liquidity. This table shows the triple-DiD estimation results with a two-month window around the event day for liquidity variables. Estimation details align with baseline specification in Table 8.

	RQS	RES	RPI	RRS	Vlm	Depth
After	3.68***	0.21***	0.35***	−0.14***	2.20***	−0.54*
	(0.52)	(0.02)	(0.04)	(0.03)	(0.25)	(0.29)
After x NasdaqStock	−0.20	0.05	0.08	−0.04	0.54	−0.27
	(0.78)	(0.05)	(0.06)	(0.04)	(0.64)	(1.03)
After x NasdaqVenue	−2.88***	−0.02***	0.04**	−0.05***	2.96***	0.71**
	(0.46)	(0.01)	(0.02)	(0.02)	(0.28)	(0.29)
After x NasdaqStock x NasdaqVenue	0.32	0.06***	0.09***	−0.03	2.76***	1.85
	(0.67)	(0.02)	(0.03)	(0.03)	(0.97)	(1.22)
R^2 (%)	0.69	1.74	1.21	0.28	1.62	0.04
N	87552	87552	87552	87552	87552	87552
Stock-Venue F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Table A8. Triple DiD Regression with a Two-Month Event Window: Trading. This table shows the triple-DiD estimation results with a two-month window around the event day for trading variables. Estimation details align with baseline specification in Table 9.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
After	−1.97***	−0.52***	−6.74***	−9.95***
	(0.22)	(0.12)	(0.63)	(0.92)
After x NasdaqStock	1.92***	0.33**	0.98	0.43
	(0.31)	(0.16)	(0.96)	(1.81)
After x NasdaqVenue	−0.45***	−0.10	−0.68	−1.03
	(0.13)	(0.08)	(0.47)	(0.71)
After x NasdaqStock x NasdaqVenue	−0.01	−0.09	3.43***	5.41***
	(0.20)	(0.11)	(0.80)	(1.65)
R^2 (%)	1.03	0.27	1.58	0.09
N	87552	87552	87551	87551
Stock-Venue F.E.	Yes	Yes	Yes	Yes

A.4 Controlling for Different Fixed Effects

In our baseline specification for the triple-DiD regression (Equation 6), we control for stock-venue fixed effects. As a robustness check, here we use an alternative specification where we control for the stock-date fixed effects:

$$
\begin{aligned}
metric_{i,t,e} = \alpha_{i,t} + \beta NasdaqVenue_e \\
+ \gamma_1 NasdaqStock_i \times NasdaqVenue_e \\
+ \gamma_2 After_t \times NasdaqVenue_e \\
+ \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue + \epsilon_{i,t,e}
\end{aligned}
\tag{7}
$$

where $metric_{i,t,e}$ is the liquidity or trading metric for stock i, traded on exchange e and on day t. $\alpha_{i,t}$ controls for the stock-day fixed effects. $After_{i,t}$ is dummy variable that equals one after October 24, 2016, and equals zero otherwise. $NasdaqStock_i$ is a dummy variable that equals one if stock i is a Nasdaq-listed stock, and equals zero otherwise. $NasdaqVenue_i$ is dummy variable that equals one if the metric is computed based on exchange e, and equals zero otherwise.

As argued in Brogaard, Ringgenberg, and Rösch (2023), the two types of fixed effects, stock-venue used in the baseline regression and stock-day used here, control for different variations. By including stock-venue fixed effects, we allow for separate intercepts for each stock-venue pair (e.g., Apple traded on Nasdaq), which is important as trading of a particular stock varies significantly across venues, especially between its listing exchange and other non-listing exchanges. For example, the trading volume of a stock is on average larger on its listing exchange (e.g., Apple is more traded on Nasdaq versus NYSE). Besides, NYSE-listed stocks are assigned to what's called DMMs (designated market makers), who have mild obligations for maintaining liquidity. In contrast, by including stock-day fixed effects, we control for unobserved effects at the stock-day level which affect trading of the stock across all venues. However, we have to impose an implicit assumption that differences within a stock across venues are constant across stocks (e.g., the effect is the same for Apple traded on Nasdaq vs. NYSE as for Nvidia traded on Nasdaq vs. NYSE).

Table A9 and Table A10 report the estimation results based on the alternative specification with stock-day fixed effects as in Equation 7 for liquidity and trading metrics respectively. It shows that coefficients on the triple interaction term are very similar to those based on the baseline

Table A9. Triple-DiD Regression with Stock-Day Fixed Effects: Liquidity. This table shows the estimation results from the triple-DiD regression

$$metric_{i,t,e} = \alpha_{i,t} + \beta NasdaqVenue_e + \gamma_1 NasdaqStock_i \times NasdaqVenue_e + \gamma_2 After_t \times NasdaqVenue_e + \gamma_3 After_t \times NasdaqStock_i \times NasdaqVenue_e + \epsilon_{i,t,e}$$

where $metric_{i,t,e}$ is the liquidity variable for stock i, traded on exchange e and on day t. $\alpha_{i,t}$ controls for the stock-day fixed effects.

	RQS	RES	RPI	RRS	Vlm	Depth
NasdaqVenue	-12.96^{***}	0.00	-0.07^{***}	0.06^{***}	5.27^{***}	10.65^{***}
	(1.04)	(0.01)	(0.02)	(0.01)	(0.34)	(0.60)
NasdaqStock x NasdaqVenue	-2.46	0.11^{***}	0.60^{***}	-0.49^{***}	9.48^{***}	16.39^{***}
	(1.66)	(0.01)	(0.04)	(0.03)	(1.50)	(2.31)
After x NasdaqVenue	-3.35^{***}	-0.03^{***}	0.01	-0.03^{**}	2.73^{***}	1.44^{***}
	(0.54)	(0.01)	(0.02)	(0.01)	(0.22)	(0.27)
After x NasdaqStock x NasdaqVenue	-0.25	0.06^{***}	0.09^{***}	-0.03	1.78^{**}	0.56
	(0.80)	(0.01)	(0.03)	(0.03)	(0.74)	(1.27)
R^2 (%)	7.06	0.49	1.93	1.38	18.60	11.42
N	171600	171597	171597	171597	171597	171600
Stock-Date F.E.	Yes	Yes	Yes	Yes	Yes	Yes

Table A10. Triple-DiD Regression with Stock-Day Fixed Effects: Trading. This table shows the triple DiD estimation results with stock-day fixed effects for trading variables. Estimation details are the same as above.

	ISOShr	OddlotShr	Cancel/Trade	Order/Trade
NasdaqVenue	1.01^{***}	2.81^{***}	0.26	3.11^{***}
	(0.11)	(0.13)	(0.52)	(0.78)
NasdaqStock x NasdaqVenue	-2.48^{***}	0.46^{**}	-11.77^{***}	-17.67^{***}
	(0.22)	(0.22)	(1.09)	(2.87)
After x NasdaqVenue	-0.39^{***}	0.42^{***}	-1.09^{***}	-1.48^{**}
	(0.12)	(0.07)	(0.37)	(0.58)
After x NasdaqStock x NasdaqVenue	0.57^{***}	-0.35^{***}	4.36^{***}	5.96^{***}
	(0.19)	(0.11)	(0.61)	(1.06)
R^2 (%)	1.01	17.28	1.45	0.13
N	171597	171597	171595	171595
Stock-Date F.E.	Yes	Yes	Yes	Yes

specification.

B Details about the Three SIP Glitch Events

Here we provide details of the three SIP glitch events used in the analysis.

Nasdaq-SIP glitch on January 3, 2013 As summarized in Table A11, on January 3, 2013, at 13:33:11 ET, the dissemination network of the Nasdaq-SIP lost connectivity, causing its even-numbered channels to cease dissemination of both trades and quote updates. A few minutes later, the remaining odd-numbered channels also ceased dissemination of trades and quote updates at 13:36:51 and 13:37:22 respectively. To present visual evidence of the glitch, Figure A2 plots the trade and quote counts from the SIP feeds and direct feeds for stocks by dissemination channel type. It shows that the number of SIP trades and quotes quickly dropped to zero for stocks in the early channels when the first glitch hit and then for stocks in the late channels when the second glitch hit. In contrast, direct feeds operated normally during the whole glitch period.

NYSE-SIP glitch on October 30, 2014 On October 30, 2014, at approximately 13:07 ET, the NYSE-SIP was hit by a hardware failure that impacted its data feed dissemination. Figure A3 in the appendix plots trade and quote counts from the SIP feeds and direct feeds for a sample of NYSE-listed stocks and Nasdaq-listed stocks respectively. It shows that at the start of the glitch, the trade and quote count of NYSE-listed stocks quickly dropped to almost zero. Note that for treated stocks, not SIP trades and quotes from all exchanges are missing. A closer examination of the SIP feeds shows that, for NYSE-listed stocks, it is either trades and quotes from Nasdaq or Bats that are missing. At about 13:34 ET, the NYSE-SIP was switched to its backup data center in Chicago.[33]

NYSE-SIP glitch on August 12, 2019 On August 12, 2019, at approximately 15:15 ET, the NYSE-SIP experienced a hardware failure of one of its network core routers, causing disruptions to the dissemination of both trades and quote updates. Figure A4 in the appendix plots trade and

[33] https://www.bloomberg.com/news/articles/2014-10-30/disaster-averted-in-nyse-stocks-as-backup-feed-kicks-in

Table A11. Channel Assignment of the Nasdaq-SIP and Outage Sequence. This table shows the symbol allocation across the six data dissemination channels of the Nasdaq-SIP. Moreover, it shows the starting and ending times of the glitch for trades and quotes in each channel.

Outage Sequence	Channel	Quote Outage Period	Trade Outage Period
"Late" Channels	Channel 1 (Symbols A-CDZ) Channel 3 (Symbols FE-LKZ) Channel 5 (Symbols PC-SPZ)	13:37:22 - 13:48:19	13:36:51 - 13:51:14
"Early" Channels	Channel 2 (Symbols CE-FDZ) Channel 4 (Symbols LL-PBZ) Channel 6 (Symbols SQ-ZZZ)	13:33:11 - 13:48:21	13:33:11 - 13:51:15

quote counts of the SIP feeds versus direct feeds for a matched sample of NYSE-listed stocks and Nasdaq-listed stocks. It shows that, during the glitch period, for NYSE-listed stocks only quote updates from NYSE Arca appear in the SIP feeds, and trades from all exchanges are missing. As the event on October 30, 2014, the operation of NYSE-SIP was later switched to the backup data center in Chicago at approximately 15:27 ET.[34].

To illustrate the three SIP technical glitch events, we plot, for each event, the number of trades and quote updates from direct feeds and consolidated feeds by the exchange. Specifically, Figure A2 plots the Nasdaq-SIP glitch event on January 3, 2013; Figure A3 plots the same metrics for the NYSE-SIP glitch event on October 30, 2014; Figure A4 plots the same metrics for the NYSE-SIP glitch event on August 22, 2019.

C Analysis of the Nasdaq-SIP Glitch Event on January 3, 2013

In addition to pooling all three SIP glitch events as above, we zoom in on the Nasdaq-SIP glitch event on January 3, 2013, and exploit its unique feature for further identification. Specifically, there are six channels through which the Nasdaq-SIP disseminates its data feeds and Nasdaq-listed stocks are allocated into the six channels *alphabetically*.[35] More importantly, as shown

[34]https://www.ctaplan.com/alerts#110000144324

[35]Ye, Yao, and Gai (2012) uses the same feature to study the potential quote stuffing behavior of HFTs and find that messages of stocks within the same channels are more correlated. The allocation of symbols in each of the six channels is according to the alphabetical order. Although the allocation rule might be such that the total message

Table A12. Summary Statistics: SIP Glitch Event on January 3, 2013. This table reports the summary statistics for all liquidity variables used in the regression for pooling all three sip glitches and for the Nasdaq-sip glitch on January 3, 2013, respectively. *RQS*, *RES* and *RRS* stands for relative quoted spread, effective spread, and realized spread respectively, and are all in basis point. *Vlm* is dollar volume in thousand dollars. *DepthNBBO* is NBBO depth in thousand dollars. *Depth5Lvl* is cumulative depth across the five best price levels of the order book. The summary statistics are computed based on time-series averages over the sample period.

Variable	N	Mean	SD	Min	50%	Max
RQS (bp)	80886	10.87	9.66	0.88	8.11	140.60
RES (bp)	21803	2.87	3.14	0.00	1.89	54.18
RRS (bp)	21803	1.61	4.15	-48.50	1.30	45.16
Vlm ($ million)	80886	8.35	49.19	0.00	0.00	2878.37
DepthNBBO ($ thousand)	80886	126.78	344.84	0.19	37.12	5985.11
Depth5Lvl ($ thousand)	80886	707.78	1796.72	11.97	196.50	23383.85

in Table A11, the Nasdaq-SIP glitch on January 3, 2013 first occurred at three even-numbered channels and occurred only a few minutes later at the other three odd-numbered channels. Thus stocks in the "Late" three channels can serve as an ideal control group during the period between the start of the first and second glitch.

Based on the same matching procedure as above, we construct a sample of 200 randomly chosen Nasdaq-listed stocks belonging to the "Early" channels, the same number of Nasdaq-listed stocks from the "Late" channels and 400 NYSE-listed stocks. In terms of the sample period, we focus on the time interval between 30 minutes before the first glitch and the start of the second glitch, i.e., between 13:03:00 and 13:36:51. Last, we run the following DiD regression:

$$metric_{i,t} = \alpha_i + \beta Period1_{i,t} + \gamma_1 Period1_{i,t} \times EarlyChannel_{i,t} + \gamma_2 Period1_{i,t} \times LateChannel_{i,t} + \epsilon_{i,t} \quad (8)$$

where $metrc_{i,t}$ is the liquidity or trading variable of stock i in time interval t. α_i is the stock fixed effects. $Period1_{i,t}$ is a dummy variable that equals one after the start of the first glitch at 13:33:11 and equals zero otherwise. $EarlyChannel_{i,t}$ is dummy variable that equals one if stock i belongs to the early channels, and equals zero otherwise. $LateChannel_{i,t}$ is a dummy variable that equals one if stock i belongs to the late channels, and equals zero otherwise. Note for the regression above, we

volume in each channel show be more or less the same so that no channel will be suffering constantly high message volume and having higher latency

Table A13. Difference-in-Difference Regression: Nasdaq-SIP Glitch on January 3, 2013. This table show regression results from the following difference-in-difference regression:

$$metric_{i,t} = \alpha_i + \beta Period1_{i,t} + \beta_1 Period1_{i,t} \times EarlyChannel_{i,t} + \beta_2 Period1_{i,t} \times LateChannel_{i,t} + \epsilon_{i,t}$$

where $metrc_{i,t}$ is the liquidity or trading metric of stock i in time interval t. α_i is the stock fixed effects. $Period1_{i,t}$ is a dummy variable that equals one between the start of the data outage at early channels and late channels, that is, between 13:33:11 and 13:36:51, and equals zero otherwise. Standard errors are clustered at the stock level. $EarlyChannel_{i,t}$ is a dummy variable that equals one if stock i belongs to one of the early channels, and equals zero otherwise. $LateChannel_{i,t}$ is dummy variable that equals one if stock i belongs to one of the late channels, and equals zero otherwise. Note that we include stocks from all early, late, and normal channels. The sample period is between 30 minutes before the outage started at early channels and the start of an outage at late channels. All variables have been defined in Table 10

	RQS	RES	RRS	Vlm	DepthNBBO	Depth5Lvl
Period1	0.04	−0.03	0.06	−1.38**	3.35	−7.89
	(0.15)	(0.05)	(0.08)	(0.67)	(4.05)	(6.14)
Period1 x EarlyChannel	0.19	0.31*	0.96***	−2.31**	−8.27	−47.03**
	(0.35)	(0.17)	(0.19)	(1.12)	(5.35)	(22.66)
Period1 x LateChannel	0.43	0.01	−0.02	0.45	−0.12	4.08
	(0.31)	(0.09)	(0.16)	(1.18)	(5.01)	(8.53)
R^2 (%)	0.09	0.08	0.31	0.04	0.05	0.61
N	80886	21803	21803	80886	80886	80886
Stock F.E.	Yes	Yes	Yes	Yes	Yes	Yes

include stocks from all channels (early, late, and normal channels). We expect γ_1 to be significant while γ_2 insignificant if the first glitch only affects stocks in the early channels. Standard errors are clustered at the stock level.

Then we zoom in on the SIP glitch event on January 3, 2013. First, Figure A5 plots several liquidity measures around the two glitches and visually shows that stocks in the early channel and late channel are affected significantly. Then we run the DiD regressions specified in Equation 8 to formally test the DiD effect. Table A13 reports the estimation results. The DiD regression is estimated over the first period of the glitch, that is when stocks of the even-numbered channels started to experience a glitch but not yet for stocks in the odd-numbered channels. The coefficient on the interaction term *Period1* × *EarlyChannel* then captures the DiD effect of the glitch-affected stocks (i.e., Nasdaq-listed stocks in the even-numbered channels) relative to unaffected stocks (i.e., both Nasdaq-listed stocks in the odd-numbered channels and NYSE-listed stocks). The results show that the market liquidity of the affected stocks worsens: relative effective spread increasing by 0.31

basis points or about 10.80% (≈0.31 ÷ 2.87) relative to its unconditional mean; cumulative depth across five best prices drops by 47.03 thousand dollars or about 6.64% relative to its unconditional mean (≈-47.03 ÷ 707.78). As in the regression where all three SIP glitch events are pooled, trading volume is significantly affected. It falls by 2.31 thousand dollars or about 27.66% (≈-2.31 ÷ 8.35) relative to the unconditional mean.

Figure A1. Cumulative Daily Changes of Liquidity and Trading Metrics for NYSE-Listed Stocks versus Nasdaq-Listed Stocks around the Nasdaq-SIP Speed Upgrade on October 24, 2016. This figure plots the cumulative daily changes in the time series of several liquidity and trading metrics for Nasdaq-listed stocks and a matched sample of NYSE-listed stocks. The vertical line represents the speed upgrade to the Nasdaq-SIP on October 24, 2016. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. The sample period is from August 29 to December 16, 2016. *RQS*, *RES*, *RPI*, and *RRS* stand for relative quoted spread, relative effective spread, relative price impact, and relative realized spread respectively, and are all in basis point. *Depth* is NBBO depth in thousand dollars. *Vlm* is trading volume in million dollars. *OddlotShr* is odd-lot trade volume as a fraction of total trade volume. *ISOShare* is trade volume via intermarket sweep order (ISO) as a fraction of total trade volume. *Cancel/Trade* is the ratio of cancel order count to total trade count. *Order/Trade* is the ratio of order volume of add order messages to total trade volume. The sample stocks consist of 296 Nasdaq-listed stocks matched with NYSE-listed stocks on price, market capitalization, trading volume, and industry. The sample period is from August 29 to December 16, 2016.

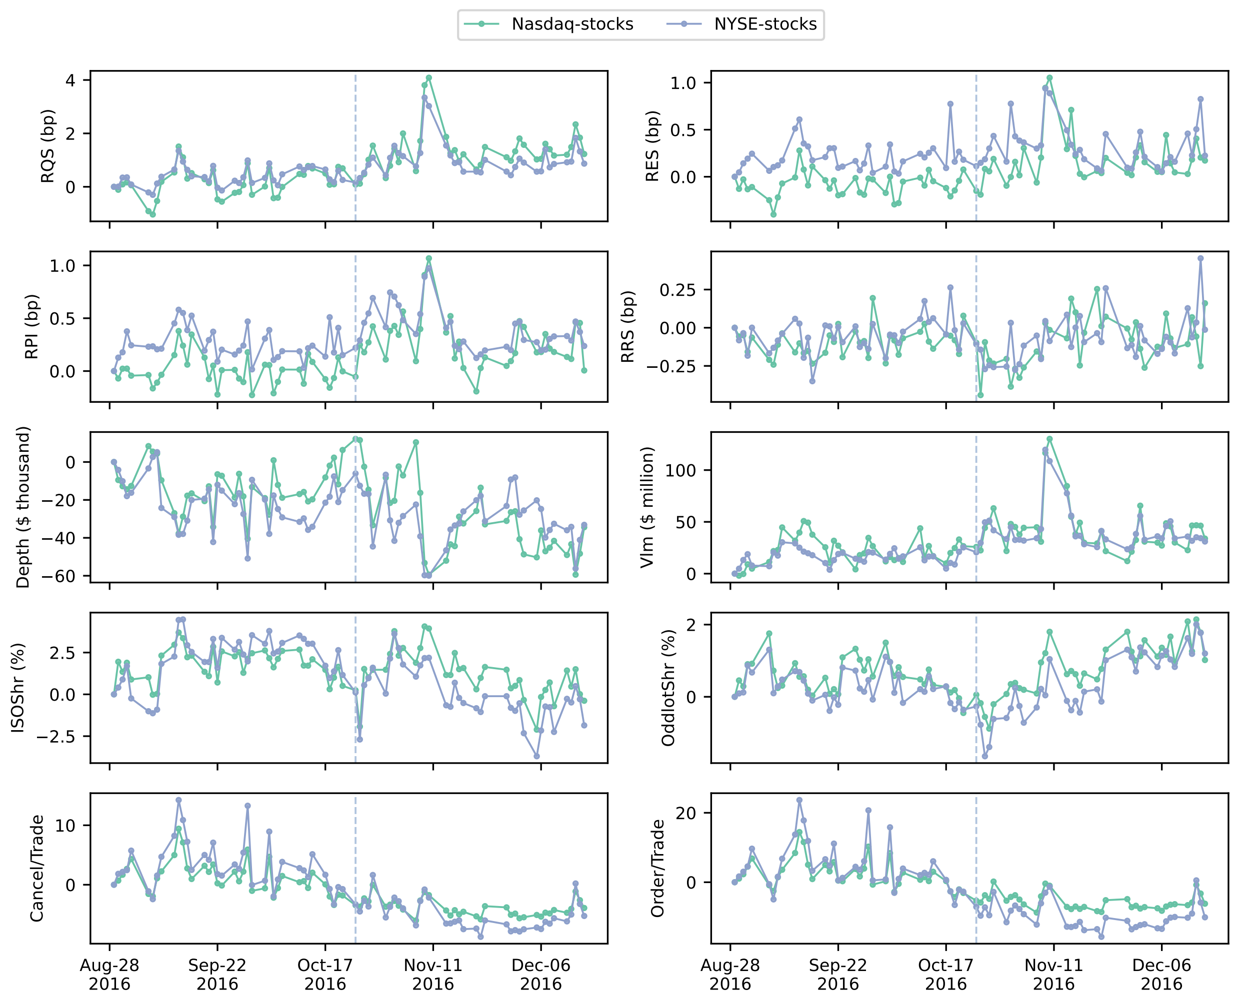

Figure A2. Nasdaq-SIP Glitch Event on January 3, 2013. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidated feeds around the Nasdaq-SIP glitch event on January 3, 2013.

(a) Number of quotes.

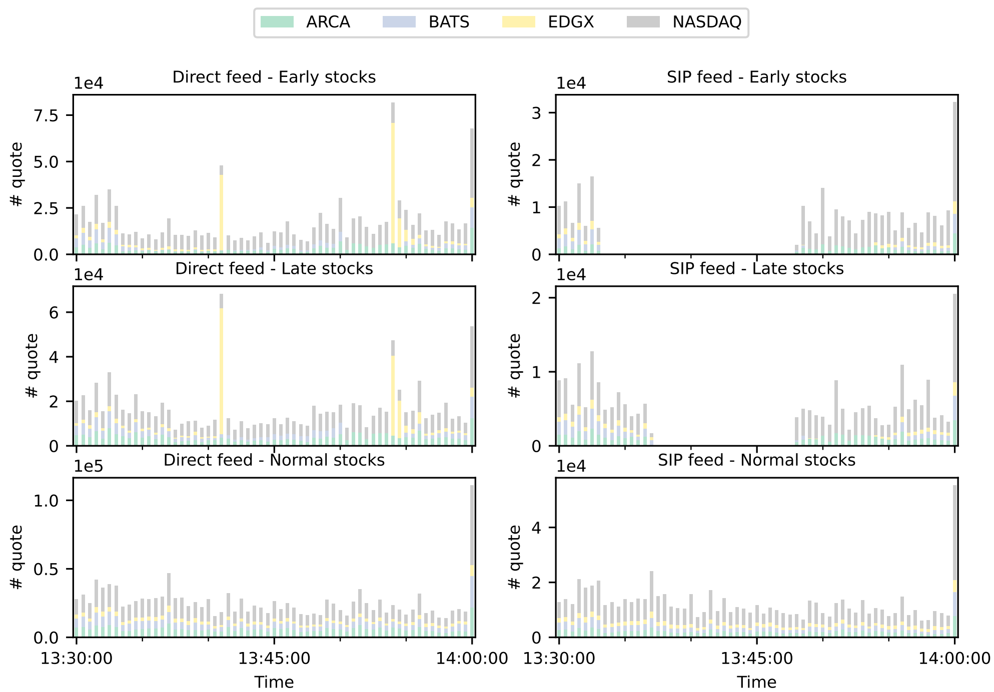

(b) Number of trades.

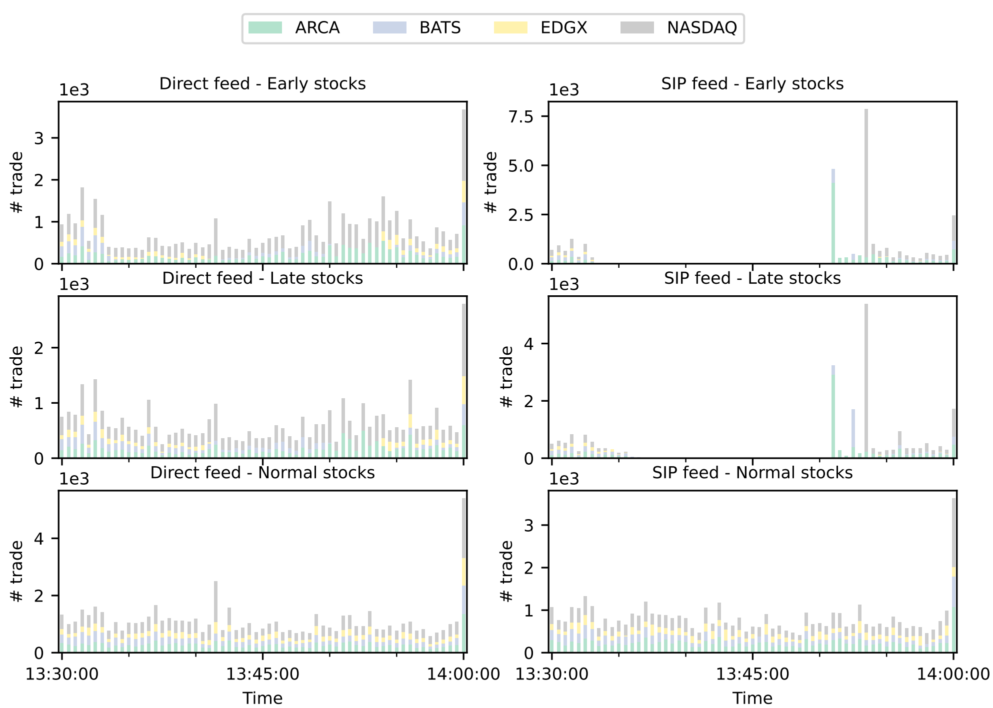

Figure A3. Nasdaq-SIP Glitch Event on October 30, 2014. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidated feeds around the Nasdaq-SIP glitch event on October 30, 2014.

(a) Number of quotes.

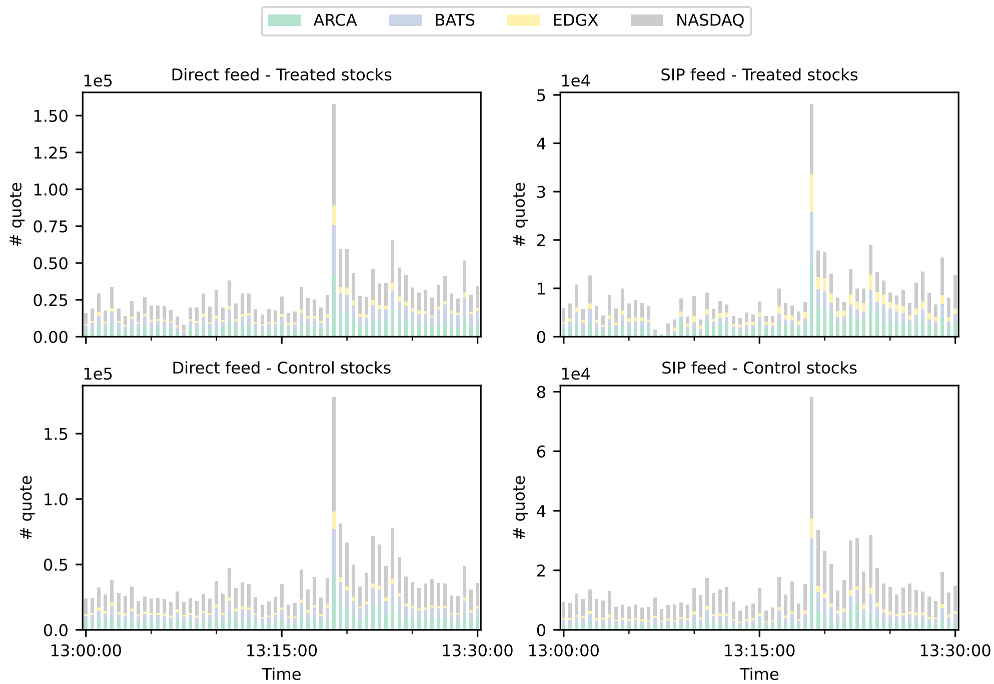

(b) Number of trades.

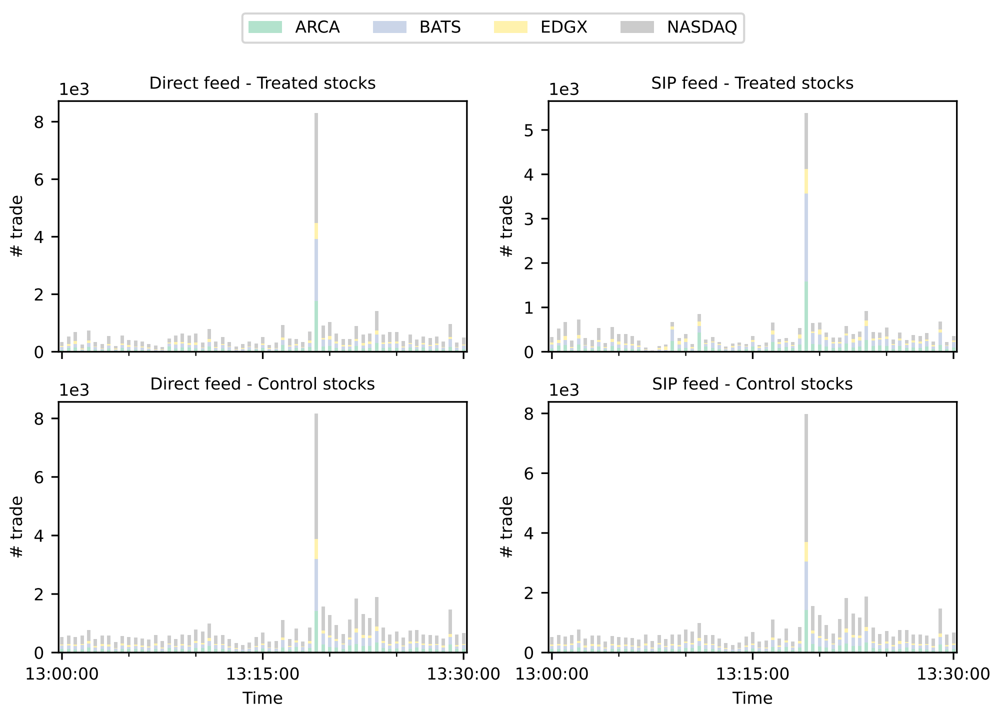

Figure A4. Nasdaq-SIP Glitch Event on August 12, 2019. This figure plots the number of trades and quote updates by exchange from direct feeds and consolidated feeds around the Nasdaq-SIP glitch event on August 12, 2019.

(a) Number of quotes.

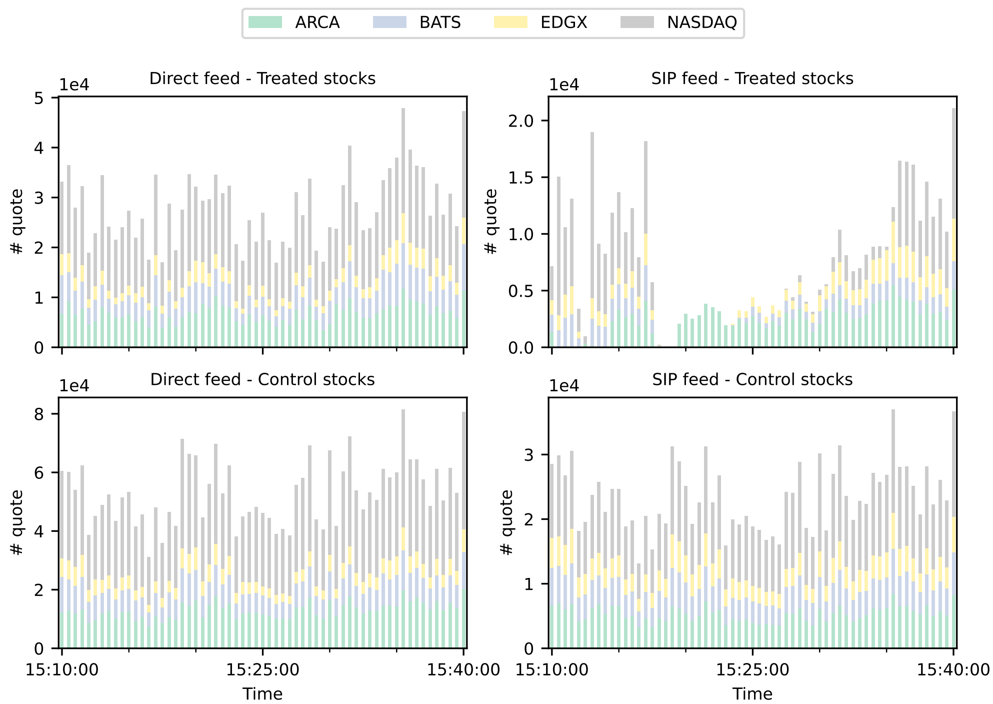

(b) Number of trades.

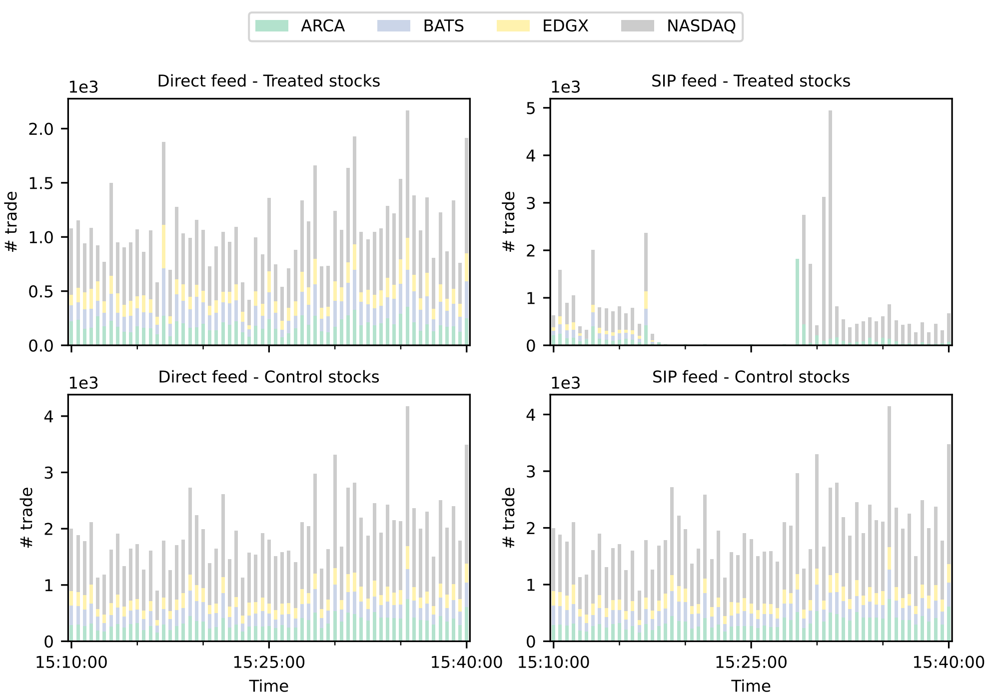

Figure A5. Liquidity Metrics around the Nasdaq-SIP Glitch Event on January 3, 2013. This figure plots several liquidity metrics for three groups of stocks respectively. The "Early" group includes Nasdaq-listed stocks belonging to the data channels that were first hit by the glitch. The "Late" groups include Nasdaq-listed stocks belonging to the data channels that were hit by the glitch later. The "Normal" group includes a matched sample of NYSE-listed stocks. *RQS*, *RES* stand for relative quoted spread and relative effective spread in basis point. *DollarVlm* is dollar trading volume in millions. *DepthNBBO* is dollar depth at NBBO in thousands. The first two vertical lines indicate the start of the glitch at the "Early" channels and "Late" channels respectively. The last vertical line indicates the end of the glitch.

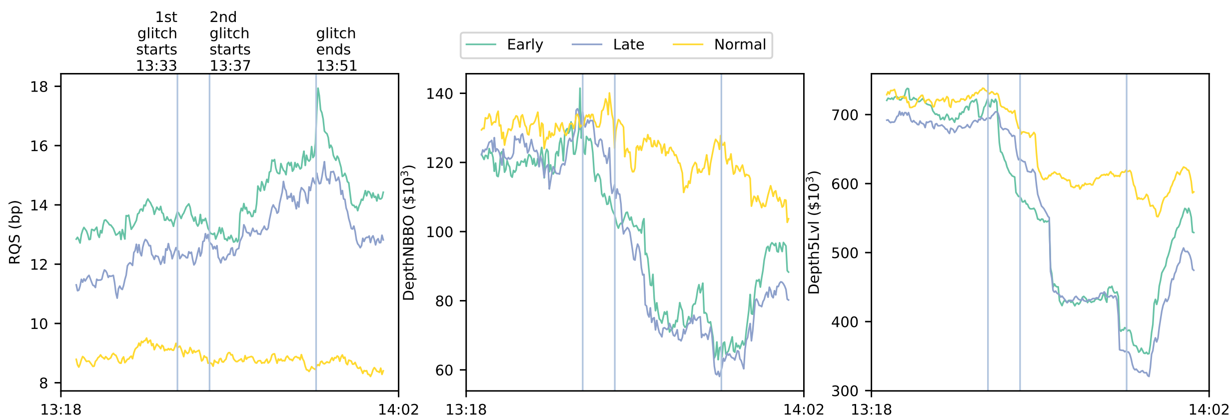

References

Aldrich, Eric M., Joseph Grundfest, and Gregory Laughlin (2017). “The Flash Crash: A New Deconstruction”. 2721922. preprint.

Aquilina, Matteo, Eric Budish, and Peter O’Neill (2021). “Quantifying the High-Frequency Trading “Arms Race””. *The Quarterly Journal of Economics* 137.1, pp. 493–564.

Bartlett, Robert P, Justin McCrary, and Maureen O’Hara (2023). “The Market Inside the Market: Odd-Lot Quotes”. *The Review of Financial Studies*, hhad074.

Bartlett, Robert P. and Justin McCrary (2019). “How Rigged Are Stock Markets? Evidence from Microsecond Timestamps”. *Journal of Financial Markets* 45, pp. 37–60.

Battalio, Robert, Shane A. Corwin, and Robert Jennings (2016). “Unrecognized Odd Lot Liquidity Supply: A Hidden Trading Cost for High Priced Stocks”. *The Journal of Trading* 12.1, pp. 35–41.

Biais, Bruno, Thierry Foucault, and Sophie Moinas (2015). “Equilibrium Fast Trading”. *Journal of Financial Economics* 116.2, pp. 292–313.

Brogaard, Jonathan and James Brugler (2020). “Competition and Exchange Data Fees”. *SSRN Electronic Journal*.

Brogaard, Jonathan, Björn Hagströmer, Lars Nordén, and Ryan Riordan (2015). “Trading Fast and Slow: Colocation and Liquidity”. *Review of Financial Studies* 28.12, pp. 3407–3443.

Brogaard, Jonathan, Matthew C. Ringgenberg, and Dominik Rösch (2023). *Does Floor Trading Matter?* URL: https://papers.ssrn.com/abstract=3609007. preprint.

Budish, Eric, Peter Cramton, and John Shim (2015). “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response*”. *The Quarterly Journal of Economics* 130.4, pp. 1547–1621.

Cespa, Giovanni and Thierry Foucault (2014). “Illiquidity Contagion and Liquidity Crashes”. *Review of Financial Studies* 27.6, pp. 1615–1660.

CFTC and SEC (2010). “Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues”.

Chakravarty, Sugato, Pankaj Jain, James Upson, and Robert Wood (2012). “Clean Sweep: Informed Trading through Intermarket Sweep Orders”. *Journal of Financial and Quantitative Analysis* 47.2, pp. 415–435.

Ding, Shengwei, John Hanna, and Terrence Hendershott (2014). “How Slow Is the NBBO? A Comparison with Direct Exchange Feeds”. *Financial Review* 49.2, pp. 313–332.

Ernst, Thomas, Jonathan Sokobin, and Chester Spatt (2021). *The Value of Off-Exchange Data*. preprint.

Foucault, Thierry, Johan Hombert, and Ioanid Roşu (2016). “News Trading and Speed: News Trading and Speed”. *The Journal of Finance* 71.1, pp. 335–382.

Hagströmer, Björn and Lars Nordén (2013). “The Diversity of High-Frequency Traders”. *Journal of Financial Markets* 16.4, pp. 741–770.

Hasbrouck, Joel (2021). “Price Discovery in High Resolution”. *Journal of Financial Econometrics* 19.3, pp. 395–430.

Hendershott, Terrence J., Marc Rysman, and Rainer Schwabe (2020). “Stock Exchanges as Platforms for Data and Trading”. *SSRN Electronic Journal*.

Hoffmann, Peter (2014). “A Dynamic Limit Order Market with Fast and Slow Traders”. *Journal of Financial Economics* 113.1, pp. 156–169.

IEX (2019). *The Cost Of Exchange Services*.

Jovanovic, Boyan and Albert J. Menkveld (2016). “Middlemen in Limit Order Markets”. 1624329. preprint.

Lee, Charles M. C. and Mark J. Ready (1991). “Inferring Trade Direction from Intraday Data”. *The Journal of Finance* 46.2, pp. 733–746.

Menkveld, Albert J. (2016). “The Economics of High-Frequency Trading: Taking Stock”. *Annual Review of Financial Economics* 8.1, pp. 1–24.

— (2018). “High-Frequency Trading as Viewed through an Electron Microscope”. *Financial Analysts Journal* 74.2, pp. 24–31.

Menkveld, Albert J. and Marius A. Zoican (2017). “Need for Speed? Exchange Latency and Liquidity”. *The Review of Financial Studies* 30.4, pp. 1188–1228.

O’Hara, Maureen, Chen Yao, and Mao Ye (2014). “What’s Not There: Odd Lots and Market Data: What’s Not There: Odd Lots and Market Data”. *The Journal of Finance* 69.5, pp. 2199–2236.

Shkilko, Andriy and Konstantin Sokolov (2020). “Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity, and Trading Costs”. *The Journal of Finance* 75.6, pp. 2899–2927.

Upson, James and Hardy Johnson (2017). “Are Odd-Lot Orders Informed?” *Financial Review* 52.1, pp. 37–67.

Ye, Mao, Chen Yao, and Jiading Gai (2012). “The Externality of High Frequency Trading”. *SSRN Electronic Journal*.