WORKING PAPER 25-01

There When You (Don't) Need It: The Reliability of Odd-Lot Liquidity

OFR
OFFICE OF FINANCIAL RESEARCH

June 5, 2025

Reed Douglas
Office of Financial Research
reed.douglas@ofr.treasury.gov

Why These Findings Are Important

Nearly half of all U.S. equity trades occur off-exchange, or without the use of an exchange. Competition from off-exchange trading venues appears to significantly affect on-exchange liquidity and prices around the time of a trade. Odd-lot orders, or those for less than 100 shares, decline in value milliseconds before an off-exchange trade occurs, only to recover a few milliseconds after. Cancellations for on-exchange orders are shown to be correlated with off-exchange trades, providing one channel through which off-exchange trading can impact on-exchange liquidity.

Key Findings

1. **Off-exchange trading activity influences liquidity, particularly for odd-lot orders and pricing.**
2. **Odd-lot liquidity increases in size but decreases in relative value in a small window of time around off-exchange trades.**
3. **Increased order cancellations may reduce odd-lot liquidity.**

How the Authors Reached These Findings

The author examines the connection between off-exchange trading and on-exchange liquidity using proprietary exchange-level data. The data include a large selection of quotes that are excluded from most databases commonly used in academic research and from the National Best Bid or Offer (NBBO). The NBBO is comprised of the best bid and best ask for a particular security, which are the aggregate of all orders in increments of 100 shares, or round lots, across exchanges. Smaller orders of less than 100 shares, or odd lots, are generally not eligible for inclusion in the NBBO.

There When You (Don't) Need It:

The Reliability of Odd-Lot Liquidity

Reed Douglas[1]

June 5, 2025

[1]U.S. Department of the Treasury, Office of Financial Research.

Abstract

Nearly half of all U.S. equity trading volume now occurs off-exchange. This study examines how off-exchange trading impacts on-exchange liquidity at high frequencies using proprietary exchange feeds. Results show that small exchange orders inside the spread, called odd lots, disappear for less than a millisecond around an off-exchange trade. This decline appears to negatively affect the prices received by many traders. On-exchange order cancellations correlate with off-exchange trades, a pattern which may propagate volatility from off-exchange venues to exchanges at extremely high speed. However, off-exchange trading seems to provide better prices than on-exchange liquidity at trade time even when compared to hidden odd-lot liquidity.

1 Introduction

U.S. equity exchanges are losing market share to off-exchange trading venues. These off-exchange venues received comparatively limited order flow 20 years ago, but now handle approximately half of all daily equity trading volume.[1] Their growth in recent years is easy to justify: They often charge lower fees than exchanges and claim to provide fast trade executions at better prices. However, they also increase the total number of trading destinations, which ultimately spreads liquidity around to different trading venues and forces traders to search for the best price across an increasingly fragmented market.[2] While investors may benefit from lower transaction fees off-exchange, they also likely face higher search costs as they compare prices across a growing set of trading venues.

The movement of trading volume off-exchange deserves attention if it negatively affects the ability of exchanges to set prices. Traditional stock exchanges still play an important role in determining the value of traded securities and disseminating that information to the broader public. Unlike their off-exchange counterparts, exchange order books are publicly observable, providing investors with near real-time information on asset values. When liquidity is directed off-exchange, effective spreads on-exchange may not be as tight, driving up trading costs for exchange-based transactions and influencing the reported prices of securities (see the discussion by Hu & Murphy, 2024). In the extreme, this process may even cause exchange market shares to decline in a downward spiral, which would limit the available pricing information.

However, it remains unclear exactly how the growth in off-exchange trading venues impacts the liquidity that is available to investors in the aggregate and on exchanges. On the one hand, off-exchange trading venues provide new sources of liquidity and are usually required to provide prices at least as good as prices observed on exchanges. This suggests that off-exchange venues should increase the total amount of liquidity available to investors. On the other hand, if enough dealers choose to provide liquidity off- instead of on-exchange, the competitiveness of on-exchange liquidity could decline. As discussed by Van Kervel & Yueshen (2023), this process could affect the quality

[1] `https://www.sifma.org/wp-content/uploads/2021/09/SIFMA-Insights-Analyzing-Off-Exchange-Trading-09-2021.pdf`

[2] The searching involved are not necessarily assuaged by the National Best Bid or Offer (NBBO). There are significant inside markets for equity orders which provide high levels of price improvement relative to NBBO. (O'Hara *et al.* , 2014)

of liquidity on exchanges.

This paper examines two questions resulting from growing off-exchange trading. First, how does the growth in off-exchange trading affect liquidity on exchanges? Specifically, and at very high frequency, it is unclear whether off-exchange trading venues provide additional sources of liquidity to investors or if they instead absorb liquidity from exchanges. Second, how well do off-exchange trading venues perform when compared to exchanges' liquidity? Studies such as Schwarz *et al.* (2023) have examined parts of this question in detail but generally rely on publicly available data sources for quotes. Instead, this study examines the quality of off-exchange trade executions relative to a type of hidden quotations that the public is generally not able to view. These quotations are called odd-lots, which are small orders that are not broadly dissemination because of their size. These quotes, however, represent significant value to investors (O'Hara *et al.* , 2014; Bartlett III, 2021; Bartlett *et al.* , 2022). This study compares the prices received off-exchange to the hidden odd-lot orders on-exchange.

The results show that the total amount of on-exchange odd-lot liquidity significantly increases milliseconds before a trade occurs off-exchange and reverts to normal levels milliseconds later. However, the relative value of odd-lot liquidity declines around a trade because the increase in odd lots occurs primarily on the same side of the order book as the off-exchange trade. The exchange odd-lot liquidity that could have matched with a trade declines in value about one millisecond before the trade occurs. These patterns correlate with on-exchange order cancellations, which increase about 5 to 10 milliseconds prior to an off-exchange trade. It appears that the most valuable exchange odd-lots have the ability to disappear a fraction of a second before they are needed. These disappearing quotes may be replaced by less-valuable ones, which could affect the quality of order execution. Despite these findings, off-exchange trading still seems to meaningfully outperform the prices available on-exchange at the time of a trade. While competition from off-exchange trading venues significantly impacts on-exchange liquidity and price setting around a trade, it provides a valuable alternative source of liquidity at the time of a trade.

To examine the connection between off-exchange trading and on-exchange liquidity, this study uses exchange-level data feeds. These data include a large selection of odd-lot quotes that are hidden from the National Best Bid or Offer (NBBO) and most standard academic databases such as Trades and Quotes (TAQ). The NBBO is the ubiquitous and easy-to-access feed of exchange order book information available to the public. Ideally, it constitutes the best bid and best ask for a particular security at any given point in time. The NBBO is broadly distributed via financial news websites and brokerage platforms, and is often available to investors at little or no cost. It aggregates all orders for increments of 100 shares, called round-lots, across exchanges. The resulting best bid and best ask, each of which can come from a different exchange, constitutes the NBBO. Odd-lot orders for less than 100 shares are generally not eligible for inclusion in the NBBO and are also missing from TAQ.

Nonetheless, odd-lots benefit investors because they can sit inside of the NBBO spread, meaning that the NBBO's 'best' bid and offer can be improved upon. The value of this odd-lot liquidity has increased significantly in recent data (see Bartlett *et al.* , 2022). In addition, the Security and Exchange Commission (SEC) Market Information Data Analytics System (MIDAS)[3] shows that odd-lot trades now represent more than half of all trades for many U.S. equity exchanges. Further, among the most widely traded stocks examined in this study, 68% of off-exchange trades were also for less than 100 shares. Therefore, on-exchange, odd-lot liquidity – not necessarily the NBBO – is the right benchmark for many trades both on- and off-exchanges. However, gaining access to this data is difficult and often requires costly direct feeds to individual exchanges, which contain information that is not publicized by NBBO feeds.

The nature of this data makes it difficult for the average investor or academic to evaluate the quality of prices received off-exchange even as an increasing number of trades find their way to these venues. This is compounded by the growth in odd lots and other forms of hidden liquidity. Even if an investor wanted to compare the price received off-exchange to the prices available on-exchange, there is a good chance the counterfactual price they would have received was hidden from the NBBO and TAQ.

To gain access to these odd lots, this study uses LSEG's Thesys and Workbench database (Thesys). This database contains a historical record of as many as 18 on-exchange data feeds, including messages to add liquidity (effectively a limit order), modify or cancel an order, and trade records. This study uses only a small subset of the full dataset that the Office of Financial Research

[3] https://www.sec.gov/marketstructure/datavis/ma_exchange_oddlotrate.html

has access to, which runs from approximately 2009 to present. Critical to this study, the Thesys data include the odd-lot quotes that would normally be excluded from the TAQ database. The Trade Reporting Facility (TRF) from the Financial Industry Regulatory Authority (FINRA) is the de-facto record of off-exchange transactions. It is not technically an exchange but is included in this data.

Two days of data are pulled for a sample of equities which successfully match the symbols used by the Thesys database. The sample equities begin as the 100 most traded stocks by transaction value between the 2020 and 2023 calendar years from a standard TAQ database. After cleaning and assigning individual trades with buy or sell indicators, the resulting sample includes about 10 million off-exchange trades. For each of these off-exchange trades, the on-exchange order book is then collected at the time the trade occurred as well as at 30 distinct points in time over a 10 second window around the trade.

The first examination of this data looks for changes in exchange odd-lot liquidity around an off-exchange trade. The tests show odd-lot liquidity nearly doubles about one millisecond before an off-exchange trade occurs, reversing and reverting to normal levels about one millisecond after.[4] This pattern is observed even when seemingly isolated and uncorrelated trades occur during periods of low trading activity. However, this increase in odd-lot liquidity does not translate into improved execution quality. Instead, much of the increase in liquidity occurs on the same side of the order book: A buy (sell) trade off-exchange appears to be associated with a significant, but short, spike on odd-lot bids (asks) on-exchange.

Next, each off-exchange trade is counterfactually 'redirected' to the best possible priced liquidity inclusive of odd lots. The results show that on-exchange odd-lot quotes that could have potentially provided improved prices disappear in the same millisecond interval that an off-exchange trade is reported. This appears to have a statistically significant and negative effect on the average trade's counterfactual exchange price. These fluctuations in odd-lot liquidity are very precise with meaningfully higher levels of odd-lot liquidity existing on-exchanges one millisecond before the trade occurred, and about one millisecond after. The extremely tight range over which this odd-lot liquidity disappears suggests that it is likely related to the off-exchange trade, though exactly to

[4] Here, odd-lot liquidity is measured as the dollar value of all odd-lot quotes inside the NBBO spread.

what end remains unclear. Possible explanations for this pattern are explored in greater detail.

The declining value of on-exchange odd-lot liquidity near an off-exchange trade (not to be confused with the doubling size or depth of odd-lot liquidity at the time of a trade) can come from two sources. First, the pattern found could result from an increase in odd-lot trading that outpaces the incoming flow of odd-lot liquidity immediately around an off-exchange trade. This is possible and examined in this study, however, the data do not include trader identities. Without identities it is difficult to know for certain if changes in order flow are directly related to an off-exchange trade or instead just correlated trading activity. Second, the patterns found could also result from an increase in order cancellations. Institutions providing quotes to exchanges can cancel their orders before they execute, making cancellations a natural and measurable way for institutions to manage their on-exchange odd-lot liquidity. Exchange cancellation messages in Thesys are voluminous, occurring commonly around trades, and they require both cleaning and special treatment. Just under 650 million cancellations occur within the 100 stock sample used in this study after netting out duplicates. Despite being common, data show a statistically significant increase in the number of on-exchange order cancellations occurring before a trade is recorded off-exchange. This suggests that certain market participants may advantageously change their on-exchange odd-lot liquidity moments before an off-exchange trade occurs. It is difficult to be certain which market participants are cancelling orders, but it could be related to predictable, correlated trading by off-exchange participants.

A process called price referencing may be related to the order cancellation results found in this study. Van Kervel & Yueshen (2023) show that dealers who participate on- and off-exchange face a conflict that can encourage them to provide worse quotes on-exchange. This conflict arises in large part from the SEC's Order Protection Rule (SEC Regulation NMS, Rule 611), which requires that off-exchange transactions, including those examined in this study, must execute at a price at least as good as prevailing quotes on-exchange. Since exchanges provide quotations in real time, this requires dealers that facilitate transactions off-exchange to check, or 'reference,' on-exchange quotes before providing a price off-exchange. Their work has shown this process can lower the competitiveness of exchange liquidity. Similar arguments are made by Hu & Murphy (2024), who add that the concentration between stocks and market-makers also appears to influence market conditions.

This study uses direct exchange feeds to test similar strategies empirically, but does not find clear evidence of price referencing at high frequency. Instead, there seems to be an incentive for off-exchange dealers to use on-exchange markets to offset trades, likely after the trade occurs. However, this study also cannot rule out this price referencing altogether even with proprietary exchange feeds because the market participant identities are still unknown, and timestamp precision is inconsistent between venues.

The empirical findings documented in this study also contribute to a growing literature that examines the separation of retail and institutional trades by selling retail trades to off-exchange dealers. Battalio (1997) shows how internalization of orders off-exchange leads to tighter spreads, but Hu & Murphy (2024) argue that the market today, which absorbs the majority of retail orders, looks very different than it did a quarter century earlier (the changes in market structure are discussed in Mahoney & Rauterberg, 2018, Chapter 5). Today, U.S. brokers send only a small number of retail orders to the exchanges that institutional orders are usually sent. Instead, brokers usually direct retail orders off-exchange, sending them to dealers or market makers. The separation of retail and institutional trades has left primary exchanges like the NYSE and NASDAQ with a comparatively small percentage of retail orders.

Proponents of this separation argue that separating retail and institutional orders should improve execution quality for retail investors. Because their trades have lower adverse selection cost and overall correlation with the market (Battalio & Holden, 2001; Ernst *et al.* , 2022; Baldauf *et al.* , 2024), dealers should be willing to facilitate these trades at better prices. Easley *et al.* (1996); Bessembinder & Kaufman (1997) discuss how this 'cream-skimming' may improve the order execution quality received by uninformed order flow. When brokers separate retail from institutional orders, dealers can be more confident they are clearing an uninformed retail order and, therefore, should provide better pricing.

Separating retail and institutional markets can benefit retail investors in theory. However, this paper argues that it remains unclear how redirected retail orders would have counterfactually changed exchange liquidity because of the high degree of on-exchange hidden liquidity. Further, the ability to improve execution quality for retail orders relies on a separation based on the average level of retail investor's informedness. Market makers would be willing to offer better pricing to retail clients only if they were reasonably sure that the order was uninformed.

This study suggests there are, in fact, significant linkages between exchanges' odd-lot liquidity – the liquidity exchanges could provide to retail investors to improve prices – and trades that occur off-exchange. Therefore, the separation of institutional and retail orders appears to be noisy at best. Not only does the value of odd-lot liquidity on exchanges change rapidly around off-exchange trades, some of these changes on-exchange occur slightly before the off-exchange trade is reported as complete. The assumption that off-exchange transactions provide retail investors with better prices depends on reliable separation across markets, an assumption for which this study finds little support.

Finally, this paper also attempts to highlight a potential inconsistency in how the quality of off-exchange trading activity is measured. Studies generally rely on measures of price improvement that compare the price received by a trade to the NBBO. However, this measure misses the growing value of odd-lot quotation activity. For that reason, measuring the quality of off-exchange trade execution with NBBO-based measures, such as price improvement, gives off-exchange trades the benefit of on-exchange odd-lot quotation activity. This study shows that odd-lots are important to off-exchange dealers' pricing strategy, and, even then, these dealers still appear to improve upon odd-lot prices at the time of a trade. This suggests that the value of uninformed traders to dealers stems from both an information component as well as the prevailing odd-lot liquidity at the time of their trade.

This paper proceeds as follows. Section 2 provides a broad overview of market structure and background information that is helpful for the remainder of this study. Section 3 covers the data used, and Section 4 contains the primary empirical results. Section 5 provides high-level commentary on how this study contributes to current regulatory and rule-writing initiatives, and Section 6 concludes.

2 Overview of Market Structure and Background

Secondary market structure in the U.S. consists of as many as 24 competing exchanges.[5] Many of these exchanges transact the same secondary market securities throughout the day and compete to facilitate transactions in the same stocks. Exchanges compete for volume with different fee structures (for example, normal versus inverted) and, ideally, by providing the best prices along with many other characteristics. Battalio *et al.* (2016) show exchange fee structures can impact where and how trades are executed.

The SEC's Regulation National Market System (Reg NMS) attempts to ensure that incoming orders are routed primarily to the trading venue with the best available price. This happens by aggregating round-lot transactions across exchanges to set the NBBO. The NBBO is the required lowest bid and highest ask that brokers must execute marketable orders at provided there is sufficient liquidity at the NBBO. This requirement means, ideally, that the NBBO is approximately the *worst* price an investor should expect to receive for a reasonably-sized order.

Trades executed without the use of a centralized exchange are still reported, usually to the FINRA TRF. While 'off-exchange trading' and 'the TRF' are frequently used interchangeably, as they will be in this paper, it is important to note that the TRF does not actually execute trades. Instead, a collection of off-exchange dealers and venues that clear trades report the resulting information about trades to the TRF. In databases such as TAQ, these trades appear to have come from a venue called TRF, but, the TRF is a record book of many underlying venues. When taken as a single entity, the value of TRF trades has grown significantly in the last twenty years.

The TRF is not a quoting exchange, so dealers that execute off-exchange must instead use a process to determine the price a trade would receive at the best-priced exchange. Such a requirement gives dealers an incentive to change their quoting procedures on-exchange if they could expect to receive an off-exchange order in the same stock for which they are providing on-exchange quotations with some probability. Dealers would want to minimize situations where they were providing the best on-exchange quote and facilitating a trade off-exchange – a situation in which they are essentially competing with themselves (see Van Kervel & Yueshen, 2023, for discussion).

[5] `https://www.sec.gov/about/divisions-offices/division-trading-markets/national-securities-exchanges`

Since the TRF does not have a quote facility, at no point does the NBBO price reflect the shares available on the TRF, and a TRF transaction necessarily does not execute against liquidity quoted in the NBBO. It is highly likely though that many firms participating on exchanges also record other transactions to the TRF.

3 Data

This study relies on a selection of data from LSEG's Thesys database and Workbench service, as well as a standard TAQ database. The Thesys database includes a historical collection of direct exchange feeds for up to 18 U.S. stock exchanges and trading venues. These direct exchange feeds include messages to add liquidity, modify orders, and cancel unfilled or partially filled orders sent to exchanges, in addition to executed trade records for both on- and off-exchange transactions.

Collecting data on all securities in Thesys would be prohibitively expensive in both time and computation. Instead, the sample is limited to the 100 stocks with the highest total dollar-value of executed trades recorded in the TAQ database between 2020 and 2023. This shifts the sample toward the most traded firms which are hopefully most representative of the average investor's trading activity. Randomly sampling firm symbols would bias the sample towards smaller and less frequently traded stocks.

Thesys is then used to reconstruct the exchange order book for all TRF trades occurring for the sample stocks for two dates outside of the ticker sample construction period, with both dates occurring in January 2024. At the time of collection, data falling after this date was limited to only about two weeks of data. The dates were randomly selected from this small sample, with the first date selected being January 16, 2024, (S&P 500: -0.4%) and the second being January 19, 2024 (S&P 500: +1.2%).[6] The data is restricted to between five minutes after market open, and five minutes before market close (9:35 a.m. to 3:55 p.m., eastern). The first and last five minutes of each trading day are excluded to avoid collecting held-over, closing and opening auction trades and quotes. All TRF trades are collected for the 100 sample securities on the two sample days which occurred at prices other than midpoint. Midpoint executions are excluded to improve the likelihood

[6] At the time of this draft, data collection was underway for the entire month of January 2024.

of successfully identifying purchases and sales using the Harris (1989) algorithm.[7] This algorithm assigns trades that take place above the midpoint as purchases and those taking place below as sells. The resulting sample includes just under 10 million TRF trades.

Thesys must be queried on specific order depths which may or may not overlap with NBBO. For example, 'Bid 1' is the highest bid at a given time for a particular security, regardless of whether it meets the requirements for NBBO. In many cases, 'Bid 1' is a small quote inside the NBBO spread, and near the midpoint, while 'Bid 2' could be lower than 'Bid 1' but still higher than the National Best Bid (NBB). The same is true for asks (equivalent to offers), except 'Ask 1' is the lowest-priced ask compared to the National Best Offer (NBO). The best five bids (Bids 1 - 5) and the best five asks (Asks 1 - 5) are collected for each trade, along with each quote's associated size.

The first and primary dataset generated is used to examine how on-exchange liquidity changes leading up to and following an off-exchange trade. This dataset brings together information about the on-exchange order book at 31 distinct time periods around each TRF trade. Specifically, the order book is collected at the following offsets around the TRF trades timestamp: plus and minus 5, 2, and 1 seconds, and 500, 400, 300, 200, 100, 50, 40, 30, 20, 10, 5 and 1 milliseconds. This grid is meant to balance the need for more granular information as the TRF trade becomes closer, with the relatively high cost of collecting this information for a large sample of trades. Collecting a constant interval (for example, every 10 milliseconds) is not feasible in this sample without significantly shortening the window because the order book must be rebuilt for each specified trade. All dates are matched with the Thesys coordinated timestamps. With five bids and five asks, along with each quotes' associated size and the NBBO bid, ask, and sizes, the sample includes $31 \times (5 + 5 + 5 + 5 + 4) = 744$ new quote data points for each of the approximately 10 million TRF trades. This dataset is identified throughout this paper as 'All Data.'

A second version of this dataset is constructed that eliminates trades for which NBBO changes for one second around the trade time (500 milliseconds on either side of the trade). As NBBO moves, so too do the prices which constitute inside-spread quotes. Therefore, variation in odd-lot liquidity is likely to occur when NBBO changes but not because of quotation activity. Eliminating

[7] Unlike trades for exchanges, TRF trades do not have a purchase or sale indicator in the Thesys database. The Harris (1989) algorithm is instead used, which cannot identify midpoint executions.

this movement helps eliminate changes in odd-lot liquidity associated with an NBBO change and not quotation activity. This dataset has significant limitations as it will bias the sample towards trades which occur during very quiet periods and also towards stocks which consistently have stable NBBO spreads that are likely correlated with price. This dataset is identified throughout as 'Fixed NBBO.'

A third version of this dataset is constructed by identifying isolated trades. A trade is isolated if it occurs with a large gap in time after the previous trade and with a large gap in time before the next trade. This dataset, therefore, eliminates periods during which large clusters of many trades arrive at once for the same stock. To identify isolated trades, the distribution of the time gaps between trades across all exchanges and the TRF are calculated for each symbol. From that distribution, the 75th percentile time gap is used as the cutoff value. All trades with at least that 75th percentile time gap between the last trade and the next trade is considered isolated. This dataset is identified throughout as 'Time Delta.'

Order cancellations and information about the original orders which were cancelled are also collected. As a reminder, the TRF is not a quoting exchange, so this dataset excludes any information about the TRF. Thesys is queried for all 'add' messages sent to the exchanges. These messages are used to submit limit and, potentially, Immediate or Cancel (IOC) orders. All modify messages, which include requests to modify an existing order to a smaller size, are also collected. Messages that modify an order to zero size are assumed to be cancels. While not explicitly flagged in the Thesys dataset, an 'add' message with an associated modify to zero with the exact same timestamp is assumed to be an IOC order. In all cases, these cancellations are collected and merged with add messages to form a dataset containing all cancelled orders and associated order information (ticker, price, size, exchange, time submitted, time cancelled, etc.). This dataset is by far the largest with more than 650 million observations over two days, but it is used in conjunction with the other datasets discussed above.

The Thesys direct exchange feeds include historical TRF trades, which should be identical to those reported in TAQ. It is important to note the TRF is not an exchange but rather a record of off-exchange transactions. Thesys includes these transactions so that everything can be queried together, but the transactions included in Thesys for the TRF are not from an exchange-based feed.

Some differences between Thesys data and those from TAQ should be highlighted. The trade data in Thesys should have the same coverage as TAQ. However, Thesys also contains a historical record of all add messages from exchanges, including odd-lot quotations sitting inside of the NBBO spread. Much of this study relies on this difference, which adds a significant number of quotations that would not otherwise be observable with TAQ. Thesys also contains a historical record of order modifications and cancellations, which will be used throughout this study. This data is also not generally available through TAQ.

Transactions in Thesys that originate from exchanges include the side of the order (buy versus sell). This information allows for an analysis of exchange data without using a signing algorithm, such as Lee & Ready (1991). Off-exchange TRF transactions, however, still do not provide this information, so the Harris (1989) method is used. While more advanced algorithms have since been developed, the Harris algorithm works very well in a sample of Thesys trades for which transaction side is included. The primary downside of the algorithm used in this study is that it necessarily removes midpoint executions. A newer algorithm proposed by Boehmer *et al.* (2021) is specifically designed to work with trades frequently reported to the TRF and appears to improve existing methods. However, their algorithm specifically targets retail investors' trades, making it less applicable for this study. Their methodology would likely exclude institutional orders, as well as retail limit orders.

3.1 Variable Construction

The data collected from Thesys are used to construct several variables. These variables are calculated separately at each offset discussed previously. The first variable is called odd-lot liquidity value, or odd-lot value, which is the dollar-denominated sum of all available odd-lots at their quote price times the number of shares. Only odd-lot quotes that are within the NBBO spread – strictly above the bid and below the ask – are used.

Next is a measure of price improvement which is taken to be the value of odd-lot liquidity relative to the actual executed trade price, P_i, of a particular trade. This measure essentially captures the total dollar value of price improvement that could have been provided by odd-lot liquidity for a trade up to the minimum size of either the TRF trade or the total odd-lot liquidity measured inside the spread. This is equivalent to a counterfactual "walking the order book."[8] This is called odd-lot price improvement (OLPI). For buys, OLPI is computed as,

$$OLPI^{B}_{i,t} = \sum_{q}^{5} (A_{q,t} - P_i) \times \min\left(S^{A}_{q,t}, S_i - \mathbb{1}_{q>1} \sum_{i=1}^{q} S^{A}_{q=i,t}\right) \tag{1}$$

where $A_{q,t}$ is the ask price at quote level q at offset t, and P_i is the actual recorded price for trade i. S_i is the size of the TRF trade, while $S^{A}_{q,t}$ is the size of the ask at quote level q at offset t. Similarly for sells,

$$OLPI^{S}_{i,t} = \sum_{q}^{5} (P_i - B_{q,t}) \times \min\left(S^{B}_{q,t}, S_i - \mathbb{1}_{q>1} \sum_{i=1}^{q} S^{B}_{q=i,t}\right) \tag{2}$$

where $B_{q,t}$ is the bid price at quote level q at offset t, and $S^{B}_{q,t}$ is the size of the bid at quote level q at offset t.

This value is meant to approximate a counterfactual measure of price improvement relative to the price recorded on the TRF. If, for example, there was significant odd-lot liquidity at the time of a small trade, this measure only considers the size of the off-exchange trade and not the larger quantity of odd-lot liquidity. Doing so provides an approximation of the price at which the off-exchange trade could have executed on-exchange had the trade been redirected.

A standard measure of price improvement is also calculated. While a different value of price improvement could be calculated for each offset separately, the price improvement at $t = 0$ will be used primarily. For buys, price improvement is defined as,

$$PI_{i,t=0} = (NBO_{t=0} - P_i) \times S_i \tag{3}$$

where $NBO_{t=0}$ is the NBBO ask price at offset time $t = 0$ such that the available shares in the quote, P_i is the actual recorded transaction price of the trade i (generally on the TRF), and S_i is the number of shares for the trade.

[8]Walking the order book refers to the situation where an order that is larger than the available liquidity first exhausts the best available quotes, before moving to the next best available quotes. This process repeats until the order is filled.

Similarly, for sells,

$$PI_{i,t=0} = (P_i - NBB_{t=0}) \times S_i \tag{4}$$

where $NBB_{t=0}$ is the NBBO bid price at offset time $t = 0$.

In most cases throughout this study, OLPI and price improvement are divided by the associated trade's total value ($price \times shares$) and expressed in basis points (bps). This improves the comparability between large and small orders. When expressed in this way, OLPI and price improvement both essentially measure rebates on the total trade's value.

To determine the side of a TRF order, which is to say whether it was a buy or sell, the Harris (1989) methodology is used. For this, the following quantity is calculated,

$$Q = 2(P_i - NBB_{t=0})/Spread + 1 \tag{5}$$

where P_i is the price of trade i, $NBB_{t=0}$ is the observed national best bid (NBB) at the time of the trade, and *Spread* is the NBBO spread at the time of the trade. If $Q > 0$, the trade was assumed to be a purchase, and if $Q < 0$, it is assumed to be a sale. This methodology excludes midpoint orders but is generally accurate when tested on a sample of trades for which the side is revealed in the Thesys data. This methodology is not necessary for exchange-based trades and quotes since Thesys provides the actual reported side of the order.

3.2 Summary Statistics

Table 1 contains summary statistics for trade-specific variables used throughout this study, measured at time $t = 0$. The sample of TRF trades contains many small trades, with a full 25% of buy (sell) orders being for one (three) shares or less and more than half of all trades being odd-lot-sized. The median buy (sell) trade value is for $886.80 ($3,757.60), though there are significantly larger trades in the data as well.

Two of the most important variables used throughout this study are standard price-improvement ('PI' in the table) and the odd-lot specific measure of OLPI. In the table, both are expressed in terms of basis points of the trade's value. We can see that the mean (median) price improvement

Table 1: Trade Summary Statistics at $t = 0$

Variable	Side	Mean	SD	10%	25%	50%	75%	90%
Shares	Buys	84.48	344.16	1.00	1.00	6.00	100.00	181.00
	Sells	118.12	417.43	1.00	3.00	25.00	100.00	233.00
Trade Value	Buys	10,620.96	25,812.19	143.02	209.95	886.80	8,484.00	29,027.32
	Sells	15,042,14	29,991.63	161.98	519.19	3,757.60	15,412.00	42,2037.00
Spread	Buys	0.05	0.13	0.01	0.01	0.02	0.04	0.12
	Sells	0.05	0.013	0.01	0.01	0.02	0.04	0.12
PI (bps.)	Buys	0.36	2.00	0.0	0.0	0.01	0.51	1.07
	Sells	0.43	3.03	0.0	0.0	0.14	0.68	1.27
OLPI (bps.)	Buys	-0.27	2.01	-1.00	-0.46	-0.00	0.00	0.00
	Sells	-0.24	2.98	-1.04	-0.51	-0.01	0.00	0.29

Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

accounts for about 0.36 - 0.43 (0.01-0.14) bps of the total trade size. At least 25% of the sample for both buys and sells occurs with PI = 0, which indicates that they occur at the NBBO quote.

OLPI is, as expected, negative on average, with a mean (median) range between -0.24 - -0.27 (-0.00 - -0.01). The negative values suggest that, on average, trades which executed on the TRF could not have performed better had they been direct to exchanges' odd-lot liquidity. If these values had been positive, it would suggest trades were consistently directed to a lower-valued venue, which does not appear to be the case. This also remains true up to at least the 75th percentile, though between the 75th and 100th percentile, there appears to be some trades which could have benefited from direct exchange routing. It is important to note, however, that extremely large or specially directed orders may be in these categories. Extremely large orders could generate a positive OLPI even though they were handled appropriately given how the measure is calculated. Table 1 measured at $t = 0$ suggests that off-exchange trading appears to provide highly competitive trade executions for investors.

Despite covering the same 100 securities for days with different average returns, there are seemingly large differences in trades depending on if they are buys or sells. This could indicate that a larger number of buys or sells were at midpoint prices, which are excluded from this study because they are difficult to properly sign. Alternatively, it could be a function of the dealer-based trading which takes place on the TRF. Since TRF trades are facilitated off-exchange and the underlying 'market' is not required to clear, we should not expect buys and sell volumes to be equal. For example, a dealer that wants to increase its net position in a stock may choose to facilitate more sell orders than buy orders. Any excess buy orders could instead be forwarded directly to a market. This could explain the seemingly significant differences between buy and sell orders.

4 Empirical Results

4.1 Do TRF Trade Prices Vary with Odd-Lot Liquidity?

Many off-exchange trades are for less than 100 shares, and therefore could have traded against odd-lot quotes had they been sent to exchanges. This is not to say that off-exchange trading does not benefit investors trading less than 100 shares. If off-exchange dealers are sensitive to or match the odd-lot prices on-exchange, then they effectively increase the total amount of inside-spread liquidity available to the market.

This study begins with a test to see if brokers and/or dealers executing trades are aware of odd-lot liquidity on exchanges and if they adjust prices accordingly. To maximize their own benefit, dealers should probably try to offer prices that are, at best, equal to either NBBO (most likely) or the prevailing odd-lot liquidity (less likely). Investors, however, benefit if dealers internalize orders off-exchange at prices better than those they could have received on-exchange. Presumably, brokers and dealers who subscribe to direct exchange feeds – which likely constitute the majority of dealers – would be able to view information about odd-lot liquidity in real time, allowing them to profit from the unobservable nature of odd lots. If they facilitated trades at NBBO but better priced odd lots were available they could capture the difference. However, the negative average OLPI in Table 1 suggests this is not what dealers do for the average trade.

Table 1 suggests that dealers are aware of odd-lot liquidity values on-exchange since OLPI is negative or zero for most trades. This means dealers are providing prices better than odd lots quotes for most trades. As a reminder, a negative (zero) value suggests that the price a TRF trade received was better than (at least as good as) the existing odd-lot liquidity on exchanges. However, establishing that TRF trades are sensitive to on-exchange odd-lot liquidity is important for determining the degree to which off-exchange trading influences on-exchange liquidity.

Table 2: Regression of Price Improvement on Relative Odd-Lot Liquidity

	All Data	Fixed NBBO	Time Delta
Rel. Odd Lot Liquidity ($, 1000s)	6.263	13.242	1.438
	(4.134)	(3.176)	(0.327)
Sell	0.169	0.088	0.164
	(6.045)	(1.904)	(3.345)
Spread	4.539	3.404	4.477
	(33.256)	(22.946)	(17.895)
Trade Value ($, 1000s)	0.001	0.001	0.001
	(18.943)	(6.682)	(20.125)
Shares (#, 100s)	-0.009	-0.010	-0.009
	(-17.318)	(-13.481)	(-18.277)
Price	4.927	0.022	6.106
	(5.013)	(2.584)	(3.009)
Symbol-Date-Minute F.E.	Yes	Yes	Yes
R^2	0.143	0.139	0.165
Obs.	10,258,563	3,425,126	2,638,205

This table contains the coefficients (t-stats.) for the regression of Price Improvement divided by trade value, expressed in basis points, on relative odd-lot liquidity. Explanatory variables include an indicator for sell orders ('Sell'; buy orders are the baseline), the NBBO spread at the time of the trade or counterfactual OLPI (Spread), the trade value calculated as the recorded trade price times the recorded trade shares, expressed in $1000s (Trade Value), the price and trade size in number of shares (Shares). Data are collected from the LSEG Thesys database covering trades on the TRF on January 16 and 19, 2024 between 9:35 am and 3:55 pm.
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

This examination begins by tabulating the value of odd-lot liquidity sitting between the NBB and NBO at the time of a trade relative to those quotes. This is equivalent to calculating OLPI for a trade size equal to all existing odd-lot liquidity. This measure calculates the dollar value of available price improvement within the spread. For example, suppose the NBB is currently \$20.00 and an inside-spread odd-lot quote to buy 50 shares is available at a price of \$20.02. Then this relative value of odd-lot liquidity for that quote is $(\$20.02 - \$20.00) \times 50 = \$1.00$. That one quote, therefore, represents \$1.00 of potential odd-lot price improvement compared to the execution price of NBBO. This value is then used to explain the trade's price improvement.

Table 2 contains the regression of price improvement on the value of relative odd-lot liquidity at the time the trade occurred ($t = 0$). The response variable in this table is price improvement on a particular trade divided by the trade's total value and expressed in basis points. Normalizing price improvement by the trade value improves comparability between large and small trades. There are three models included in the table: One for each of the three primary datasets examined.

In the models that include All Data and Fixed NBBO, the coefficient of interest is the odd-lot liquidity which is positive and statistically significant. As the dollar value of odd-lot liquidity on-exchanges increases relative to the NBBO, the prices provided by dealers internalizing orders off-exchange improves as well. This value appears to be significant: for every $1,000 of additional potential price improvement in odd-lot liquidity, dealers appear to provide about 6 bps more price improvement for the full data set, and about 13 bps when NBBO remains fixed. The coefficient for the Time Delta dataset is statistically insignificant. These results compare to an average of about 0.4 bps in actual price improvement received by trades in the sample.

These findings suggests that dealers operating off-exchange are likely aware of the value of hidden on-exchange odd-lot liquidity, as expected. This is reassuring, but it is by no means given. A dealer that is unaware of on-exchange odd-lot liquidity could generate more profit from trade execution had they priced all trades at NBBO minus some small, sub-penny premium. For example, a dealer receiving a client's buy (sell) order could have provided the client the NBO (NBB) minus (plus) some small sub-penny amount. Such a policy would have still provided the retail client with a price-improved trade, did not violate the NBBO, and would have nearly maximized the dealer's profits.

That, however, does not appear to be the strategy employed by dealers reporting to the TRF. Instead, these dealers appear aware of the value of odd-lot liquidity on exchanges and price their off-exchange trades accordingly. This result is somewhat surprising because investors, particularly retail investors whose trades are almost always directed to the TRF, would likely be unaware if their trade had not received odd-lot prices. This is due to the hidden nature of odd-lot quotes. Therefore, it seems that dealers provide price-improved trades even when the investor placing the trade would most likely not be able to monitor odd-lot prices.

The result that off-exchange prices are sensitive to on-exchange odd-lot liquidity likely benefits investors. It implies that the actual amount of available inside-spread liquidity is larger than what exchanges report. This is because off-exchange dealers are willing to provide liquidity at inside-spread prices not included in exchange's odd-lot liquidity. This may increase the total amount of valuable inside-spread liquidity available to the market, improving the chances that investors can execute both on- and off-exchange at improved prices.

4.2 Odd-Lot Liquidity around Trades

The previous section examined odd-lot liquidity at the time of a TRF trade ($t = 0$). This section examines how odd-lot liquidity varies *around* TRF trades ($t < 0$ and $t > 0$). Two variables are used to measure the value of on-exchange odd-lot quotes. The first is simply the total value of on-exchange odd-lot liquidity that is inside of the SIP-disseminated spread. The second is a relative measure of the odd-lot liquidity value and compares the prices that off-exchange TRF trades actually received to the prices they could have received from the best available on-exchange odd-lot quotes. This value is called OLPI and is measured at 30 distinct time periods around each TRF trade (see Section 3.1). OLPI is essentially the best possible dollar profit an investor could have received had their trade been sent to the exchange or exchanges with the best available odd-lot liquidity instead of the TRF. OLPI is divided by the trade's total value and expressed in basis points to make it comparable between both large and small trades.

The 30 non-contemporaneous offsets examined are evenly split pre- and post-trade. When compared to the contemporaneous data ($t = 0$), these offsets can provide useful information about how the order book changes before and after a trade if at all. If there is no change in the odd-lot liquidity value or OLPI in a window around the trade at $t = 0$, then the execution of an off-exchange trade may be unrelated to liquidity on-exchange. Such a finding would be consistent with the TRF providing additional liquidity to the market without impacting existing on-exchange liquidity. If instead the pre- and post-trade offsets have significantly different levels of odd-lot liquidity and OLPI than at $t = 0$, then TRF liquidity may not be an additional source of liquidity but instead, an alternative source.

As a reminder, trades reported to the TRF are negotiated off-exchange and do not execute against standing limit orders on exchanges. Because of this, a trade reported to the TRF is necessarily not paired with on-exchange odd-lot liquidity. Therefore, there should be no mechanical decline in the availability of odd-lot liquidity immediately surrounding an off-exchange trade.

Figure 1 contains two figures showing the average odd-lot liquidity and OLPI for the sample of

Figure 1: Odd-Lot Liquidity and OLPI, All Data

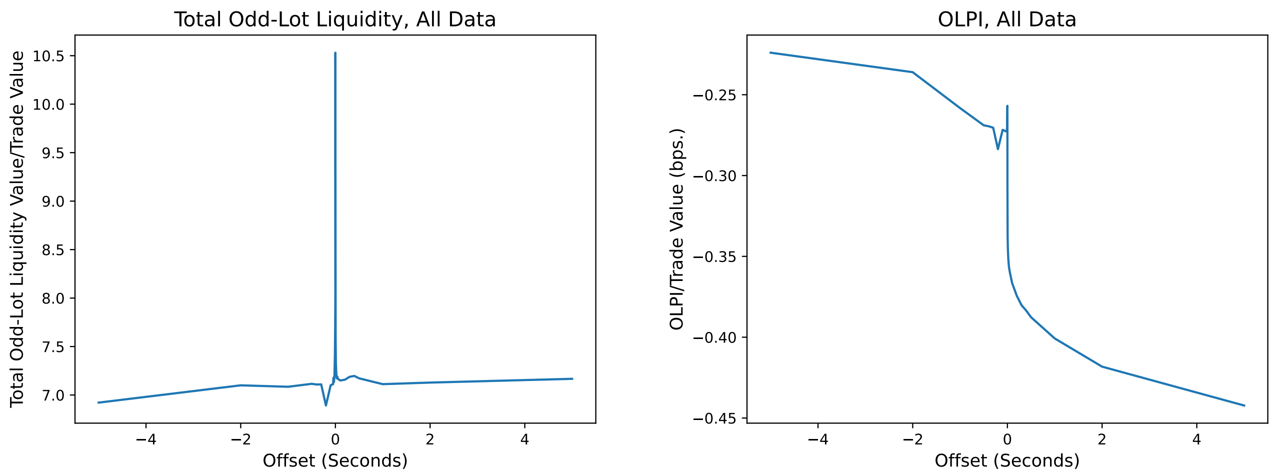

(a) Total Odd-Lot Liquidity/Trade Value, All Data

(b) OLPI/Trade Value, All Data

OLPI/Trade Value is the odd-lot price improvement for a particular trade divided by the trade's value expressed in basis-points (bps.). Total Odd-Lot Liquidity is the total value of inside-spread odd-lot liquidity. Total Odd-Lot Liquidity/Trade Value is a ratio of the odd-lot liquidity level at the time of that trade relative to the trade's value; a value of 1 means that the odd-lot liquidity at that time was equal to the size of the trade's value. These graphs contain all data in the full sample (All Data).
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

all trades. Figure 1a contains the average odd-lot liquidity around the average trade. It is expressed as a ratio: A value of 2, for example, indicates that at the time of the trade, there was enough odd-lot liquidity on exchanges to absorb the average trade twice over. The x-axis is the time offset used when collecting the exchanges' order books, with $t = 0$ being the time the trade was actually recorded. Negative offsets represent the order book seconds before the trade time, while positive offsets are seconds after.

Figure 1a shows that there is enough odd-lot liquidity to clear the average TRF trade about seven times over. This ratio increases to more than ten at the time of the trade before quickly returning to nearly the same pre-trade level of about seven. The large spike in odd-lot liquidity at the time of a TRF trade suggests that the amount of odd-lot liquidity inside the spread at the time a trade occurs is not the average amount during the trading day. In fact, if one were to measure odd-lot liquidity using only the contemporaneous order book, they would assume that the average available odd-lot liquidity was about 50% larger (10.5x) than it is during the majority of the trading day (7x). These results are surprising: If price referencing activity was taking place in odd-lots at high frequency, we should expect a decline and not an increase at the time of the trade. There is a clear benefit to dealers when facing lower on-exchange odd-lot liquidity around TRF trades because they could set prices without the pressure to compete with on-exchange odd-lot liquidity. However, these findings suggest that odd-lot liquidity rapidly increases immediately around an off-exchange trade, which would seemingly increase the competition these dealers face.

As the amount of available odd-lot liquidity increases significantly around a TRF trade, Figure 1b shows that the relative value of that liquidity to the investor declines. OLPI/Trade Value declines slowly as event time approaches $t = 0$, before dropping significantly within a millisecond or two of the actual trade time. The decline in OLPI slows but continues decreasing for several seconds after the trade occurred. The y-axis is always negative, which means that odd-lots on exchanges could not have improved the prices received by investors at any of the points in time examined. At the time of the trade ($t = 0$), the OLPI was approximately -0.26 bps, suggesting that the TRF provided better prices than what could have been expected from on-exchange odd-lot liquidity.

The shape that is generated in Figure 1b is not the expected pattern if TRF trading simply provided additional liquidity to the market. Instead, the pattern suggests that off-exchange trades appear to coincide with significant declines in the relative value of on-exchange odd-lot liquidity, and those declines are persistent for up to several seconds after the trade occurs.

These findings may not be surprising if the TRF was a bespoke exchange. For example, if a broker or trader were to consistently route orders to the exchange with the highest levels of odd-lot liquidity, then we should expect, on average, that each exchange would generate a negative OLPI at $t = 0$. This is by virtue of the way OLPI is calculated. However, the TRF is not an exchange, and instead, many of the trades routed to dealers and reported to the TRF avoid the exchange routing process. It is therefore difficult to reconcile how trades which avoid the exchange-routing process still appear to have been consistently routed to the best venue – the TRF – and why the value of exchange's OLPI seems to change immediately around those trades.

The pattern shown in Figure 1a could occur if a significant number of trades in the same stock were to occur at the same time, with some sent to exchanges and others sent to off-exchange dealers. This would imply that for some non-structural reason, most off-exchange trades end up

Figure 2: Odd-Lot Liquidity and OLPI, Isolated Trades

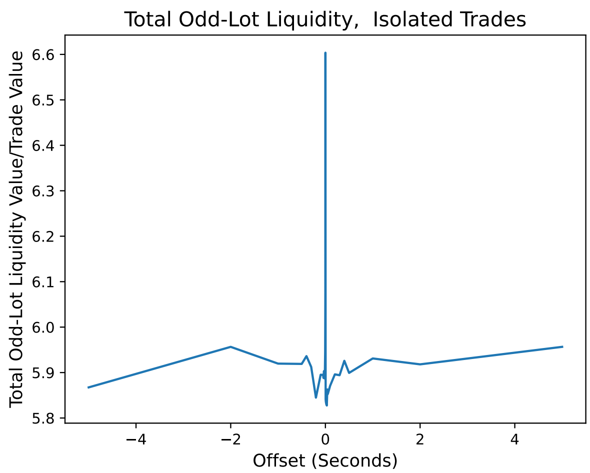

(a) Total Odd-Lot Liquidity/Trade Value, Isolated Trades

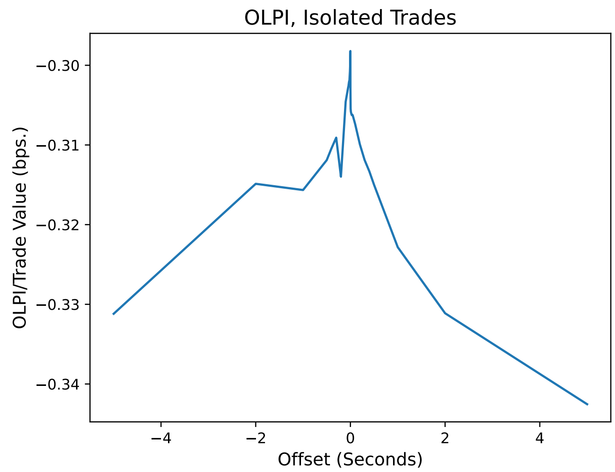

(b) OLPI/Trade Value, Isolated Trades

OLPI/Trade Value is the odd-lot price improvement for a particular trade divided by the trade's value expressed in basis-points (bps.). Total Odd-Lot Liquidity is the total value of inside-spread odd-lot liquidity. Total odd-lot value/Trade Value is a ratio of the odd-lot liquidity level at the time of that trade relative to the trade's value; a value of 1 means that the odd-lot liquidity at that time was equal to the size of the trade's value. These graphs contain all data in the Isolated Trades sample, which are trades that occurred when there was no change in the NBBO price for at least 500 milliseconds before the trade, and 500 milliseconds after.
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

trading when odd-lot liquidity is significantly higher. There may be some precedent for this pattern, albeit at a lower frequency. Correlated trading among small, likely retail, investors has been well examined, including herding patterns discussed by Barber *et al.* (2008). Kumar & Lee (2006) shows that retail traders tend to buy and sell in the same stocks, while Dorn *et al.* (2008) shows similar patterns at frequencies as high as the daily level. However, the patterns found in Figure 1 show that variation in odd-lot liquidity is very high frequency, requiring similar levels of correlated trading at the sub-second level. This would require a level of attention and access to real-time exchange data that seems unlikely for many off-exchange investors, including retail traders who frequently use the TRF. Nonetheless, the potential for correlated trading is now examined.

Figure 2 contains only those trades that are a large distance away from the previous trade and the next trade, across all exchanges for that symbol. Here, a large distance is defined as the 75th percentile of time between trades for that symbol. To arrive at this data set, each symbol's distribution of time between trades across all exchanges is calculated. Then, only those trades which have a time from the previous and next trades greater than the 75th percentile of trades are included. This means that the trades identified are significantly isolated from other trades for that symbol, even across other exchanges. This should help to significantly eliminate correlated trading within that symbol around the time of the trade.

Figure 2 shows that even the most isolated trades are subject to cross-venue influences. Despite having relatively large gaps of time between trades for that symbol, Figure 2a shows that the total value of odd-lot liquidity once again increases, moving from about 5.9 to 6.6 times the size of the average trade. This indicates that even for highly separated trades, the spike in odd-lot liquidity seems to be related to the fact that an off-exchange trade is occurring, not when it occurred (relative to other trades). However, unlike Figure 1b, Figure 2b shows that OLPI rises slowly as t approaches zero and then declines following the trade. This pattern suggests that trading occurs roughly when odd-lots on exchange are most valuable, but once that trade occurs, odd-lots begin a decline in value. This suggests that odd-lot liquidity is influenced by off-exchange trading, and that off-exchange trading may negatively impact the relative value of exchange odd-lots.

4.2.1 Regressions

Figures 1 and 2 show key patterns that are examined more closely in this section. First, odd-lot liquidity seems to increase significantly around the time of a trade but returns to its previous level shortly after the trade is complete. Second, trades reported to the TRF appear to get better prices than on-exchange liquidity. A statistical test is used to check if these changes in odd-lot liquidity and prices are statistically significant.

The first trade-level analysis is an event study of inside-spread odd-lot liquidity around the time of an off-exchange trade. The movement of liquidity or volume between on- and off-exchange venues cannot be measured directly because the data used do not include trader identities. For this reason, an event study is used to provide a statistical estimate of the relationship. The liquidity measure used includes only the first five levels of odd-lots that were inside of the SIP-reported NBBO at the time of the trade's offset.

If trading off-exchange influences liquidity on-exchange at all, it may do so in at least two different ways. An off-exchange buy or sell trade may induce either buying or selling activity on-exchange. If an off-exchange buy (sell) encourages a dealer to replenish their inventory on-exchange, this would be visible as an on-exchange buy (sell). If instead an off-exchange participant was actively trying to eliminate (purchase) shares both on- and off-exchange, potentially affecting the quality of trade execution, then we may see that an off-exchange buy (sell) is accompanied by other on-exchange sells (buys). The discrepancy is because the dealer's sell quotation on-exchange is the same type of liquidity provision as the investor's buy of the dealer's off-exchange sell. The former method would allow dealers to capture a small inside-NBBO odd-lot spread over very short periods of time, while the latter would suggest that dealers use all venues to manage their inventory.

The model used in this first event-study is the following:

$$OLL_{s,i,t} = \alpha_i + \sum_{h=-H}^{H} \beta_h D_{i,t}^h + \epsilon_{s,i,t} \tag{6}$$

where $OLL_{s,i,t}$ is the odd-lot liquidity on side s for trade i and offset t. Here, side indicates if the liquidity is matched on the same side (buys off-exchange to buys on-exchange) or if it is eligible opposite-side liquidity (buys off-exchange to sells on-exchange), and H is the set of all non-negative offsets. α_i is a trade-level fixed-effect, and $D_{i,t}^h$ equals 1 when $t - h = 0$ and is 0 otherwise. As a reminder, the on-exchange order-book is queried 31 different times for each trade, which is the variation captured by this model. There are just over 10 million fixed-effects included in the models run with all data.

The results of this event study are included in Table 3. This table contains six columns. Models (1) - (3) include the odd-lot liquidity at the time of the trade that was on the opposite side of the exchange's order book. That is, these models regress the total value of inside-NBBO spread odd-lot buys (sells) orders around off-exchange sell (buy) orders over the 30 offsets examined in this study ($t = 0$, the contemporaneous order book, is the 31st offset, but is removed as the baseline in this regression). These are labeled as 'Odd-Lot Liquidity, Eligible' because, presumably, the off-exchange order could have been match with all or part of these odd-lots had it been directed on-exchange instead of off-exchange. Models (4) - (6) include the quotes on the same side of the on-exchange order book. For off-exchange buy (sell) orders, these are the on-exchange buy (sell)

Table 3: Regression of Odd-Lot Liquidity on Offsets

	Odd-Lot Liquidity, Eligible			Odd-Lot Liquidity, Same Side		
	(1)	(2)	(3)	(4)	(5)	(6)
Offset (sec.)	All Data	Fixed NBBO	Time Delta	All Data	Fixed NBBO	Time Delta
-5	136.80	366.52	240.49	-3270.22	156.64	-515.41
	(19.07)	(64.31)	(28.74)	(-145.67)	(27.65)	(-33.45)
-2	77.77	262.10	209.52	-3230.29	107.02	-472.63
	(11.57)	(58.65)	(29.66)	(-144.97)	(23.50)	(-31.58)
-1	38.03	138.84	164.79	-3205.48	35.71	-445.36
	(5.85)	(41.55)	(27.00)	(-144.30)	(9.98)	(-30.84)
-0.5	-10.16	14.46	132.45	-3167.92	-40.69	-420.54
	(-1.55)	(7.08)	(22.03)	(-142.79)	(-17.75)	(-29.66)
-0.4	-21.71	7.05	122.98	-3163.89	-43.21	-412.85
	(-3.39)	(3.77)	(23.83)	(-142.67)	(-20.18)	(-29.26)
-0.3	-39.38	6.60	112.25	-3150.33	-39.72	-402.85
	(-6.23)	(3.84)	(21.69)	(-142.08)	(-20.30)	(-28.77)
-0.2	-66.40	5.09	84.93	-3146.70	-36.33	-384.53
	(-10.73)	(3.41)	(18.40)	(-141.34)	(-20.24)	(-27.61)
-0.1	-80.44	9.11	67.63	-3099.07	-30.77	-368.95
	(-12.96)	(7.47)	(17.98)	(-139.52)	(-20.29)	(-26.16)
-0.05	-89.30	9.59	47.89	-3081.09	-28.32	-359.68
	(-13.98)	(9.80)	(14.28)	(-139.51)	(-22.48)	(-26.46)
-0.04	-96.71	9.71	36.38	-3069.26	-27.39	-359.44
	(-15.93)	(10.58)	(11.63)	(-138.29)	(-22.79)	(-26.44)
-0.03	-96.78	10.05	33.05	-3046.03	-26.03	-356.68
	(-15.91)	(11.79)	(10.64)	(-137.25)	(-22.64)	(-26.28)
-0.02	-90.88	10.01	20.40	-3012.72	-24.46	-359.97
	(-14.38)	(13.26)	(7.88)	(-135.21)	(-22.32)	(-26.73)
-0.01	-52.13	8.63	17.98	-2956.50	-21.21	-351.62
	(-7.79)	(13.19)	(6.23)	(-132.48)	(-20.34)	(-26.17)
-0.005	12.53	8.13	25.39	-2872.93	-18.10	-341.64
	(1.56)	(12.42)	(7.90)	(-128.44)	(-18.76)	(-25.37)
-0.001	244.60	4.52	37.52	-2627.45	-13.54	-331.26
	(26.07)	(11.56)	(7.70)	(-119.52)	(-18.34)	(-24.57)
+0.001	182.31	0.23	32.35	-1501.84	20.24	-189.37
	(28.12)	(0.23)	(11.44)	(-57.09)	(14.16)	(-10.32)
+0.005	109.89	-10.88	32.91	-2683.21	2.26	-408.26
	(14.54)	(-16.37)	(13.86)	(-121.79)	(2.29)	(-30.28)
+0.01	44.06	-12.81	31.35	-2848.39	-3.40	-415.82
	(6.85)	(-18.34)	(12.26)	(-128.84)	(-3.32)	(-30.82)
+0.02	-12.77	-14.01	29.06	-2963.88	-4.91	-419.81
	(-2.14)	(-18.34)	(10.76)	(-134.36)	(-4.54)	(-30.82)
+0.03	-19.40	-14.17	33.30	-2990.86	-4.62	-414.63
	(-3.24)	(-16.48)	(12.84)	(-135.55)	(-3.99)	(-30.45)
+0.04	-7.08	-13.50	42.41	-2996.90	-5.21	-402.73
	(-1.09)	(-14.34)	(12.92)	(-135.68)	(-4.28)	(-29.14)
+0.05	-19.07	-13.13	50.68	-3004.04	-3.67	-396.57
	(-3.22)	(-13.30)	(12.93)	(-136.24)	(-2.80)	(-29.18)
+0.1	-21.98	-9.70	71.59	-3062.60	-1.77	-393.19
	(-3.59)	(-7.85)	(16.80)	(-138.45)	(-1.17)	(-28.54)
+0.2	-1.64	-0.59	93.16	-3098.33	5.14	-405.58
	(-0.27)	(-0.39)	(23.21)	(-140.03)	(2.61)	(-29.10)
+0.3	18.90	9.09	115.77	-3119.92	8.26	-417.63
	(3.06)	(5.38)	(26.20)	(-140.98)	(4.23)	(-29.69)
+0.4	36.90	20.19	130.04	-3145.17	14.42	-415.14
	(5.89)	(10.99)	(27.31)	(-141.89)	(6.88)	(-28.77)
+0.5	53.74	38.85	144.62	-3163.78	29.10	-432.64
	(8.47)	(18.58)	(28.44)	(-142.74)	(10.09)	(-30.54)
+1	99.29	169.70	182.35	-3205.49	102.32	-450.10
	(15.13)	(48.33)	(30.36)	(-144.16)	(22.15)	(-30.90)
+2	138.29	278.34	208.97	-3233.18	121.26	-481.37
	(20.42)	(56.47)	(29.45)	(-145.18)	(25.78)	(-32.51)
+5	208.58	393.72	260.90	-3248.37	170.07	-510.44
	(29.42)	(65.79)	(32.30)	(-145.24)	(29.47)	(-33.45)
R^2	0.46	0.87	0.75	0.30	0.87	0.66
Obs.	315,565,043	105,282,791	81,527,288	315,565,043	105,282,791	81,527,288
Trade F.E.	Yes	Yes	Yes	Yes	Yes	Yes

This table contains the coefficients (t-stats.) for the regression of counterfactual OLPI on time offsets. Data are collected from the LSEG Thesys database covering trades on the TRF on January 16th and 19th, 2024 between 9:35 am and 3:55 pm. Offsets are seconds before (negative) or after (positive) the time that the trade occurred (offset of 0.0).
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

odd-lots. The off-exchange order could not have executed against these because they are on the wrong side of the book. However, this liquidity is examined because it could indicate that dealers participating off-exchange are effectively reversing the trade and offloading it onto an exchange using odd-lot liquidity.

Starting with Model (1), which contains all data in the sample and uses trade-level fixed-effects, we can see that odd-lot liquidity is negative and statistically significant between 400 and 10 milliseconds before a trade occurs. Then, in the five or so milliseconds before the trade occurs, odd-lot liquidity increases and is higher than at the time of the trade. One millisecond before the trade, for example, on-exchange, odd-lot liquidity was about $245 higher on average than it was at the time of the trade. One millisecond after the trade, the same liquidity was $182 higher on average, and the coefficients remain positive and statistically significant between 10 and 20 milliseconds after the trade. The coefficients then become negative between 20 and about 200 milliseconds after the trade. For the furthest intervals measured, 5, 2 and 1 seconds before and after the trade occurred, on-exchange odd-lot liquidity is higher on average, with positive and statistically significant coefficients for all of these offsets. The model has an R^2 of 46%.

The findings in Model (1) suggest that the odd-lot liquidity available on-exchanges is likely impacted by off-exchange trading. This analysis relies on the tight window in which this change occurs around an off-exchange trade. In the millisecond before and after a trade, odd-lot liquidity drops by about $245 on average and then nearly recovers a millisecond after the trade. It is hard to imagine how this could be related to something other than the off-exchange trade. Correlated trading could cause this pattern, but it is difficult to identify entities that would find this activity beneficial and are capable of trading over such a tight window. Retail investors, for example, are known to trade in a correlated fashion, particularly when using the internet to discuss trading ideas. However, it seems unlikely these investors would have the ability to time trades to within 1 millisecond of each other. Further, these intervals are tighter than the reported execution times for many retail brokers. Professional institutional traders could time trades with high precision, but it is hard to understand what benefit they would gain by trading on the TRF and exchanges simultaneously.

Next, a sample of trades for which the NBBO did not change are examined. As a reminder, these are trades for which the NBB and NBO did not change over the offsets in the sample. Relatively few trades have NBB and NBO that remain fixed for ten seconds, so the filter is applied only to those times between the -0.5 and +0.5 offsets, inclusive.[9]

Model (2) examines only those trades in the Fixed NBBO sample. The patterns in Model (2) are more consistent pre-trade than those observed in Model (2). All offsets occurring before the $t = 0$ have positive and statistically significant coefficients, suggesting that odd-lot liquidity was usually higher before the trade than at the time of the trade. Similar to Model (1), Model (2) indicates that on-exchange odd-lot liquidity declines significantly at the time of a trade. One millisecond prior to the off-exchange trade, on-exchange odd-lot liquidity was an average of \$4.52 higher and was roughly the same level one millisecond after the trade. Approximately 5 milliseconds after the trade, coefficients become negative (-\$10.88), meaning that the value of liquidity is lower than at the time of the trade. Negative coefficients continue until about 200 milliseconds after the trade, where they switch again to positive. The model generates an R^2 of 87% with trade-level fixed-effects.

Model (2) once again suggests that on-exchange odd-lot liquidity is impacted by off-exchange trading. The pattern in this model suggests that off-exchange trading appears to remove on-exchange odd-lot liquidity and that liquidity may take a while to recover, remaining insignificant or negative for approximately 200 milliseconds. As discussed above, these patterns could occur as a result of changes in the quotation activity or because of correlated trading on-exchange. Here, correlated trading simply means that multiple trades are sent to multiple venues including the TRF very near one another. Correlated trading could indicate that, within the same millisecond, a market participant is effectively capturing on-exchange liquidity as a result of the off-exchange trading. Identifying this activity for certain would require market participant identifiers, but those identifiers are not included in the data used in this study. However, Model (3) attempts to control for highly correlated trading activity across market centers.

The last sample examined are all trades that are seemingly isolated from other trades such that for each symbol, the trades occurred with a significant gap between that trade, the previous trade,

[9]One limitation of this analysis is the possibility that NBB and/or NBO could vary between offsets, but, by chance, return to the same level at each of the sampled offsets. The nature and size of this data means that the quotes are pulled as snapshots, and it cannot be known for certain that the NBB or NBO did not vary between those snapshots. However, this would require both quotes to return to the same level at 26 distinct periods in a 1 second interval. While unlikely, it is possible for this to occur.

and the next trade. The significant gap is taken to be the 75th percentile time between trades within each symbol over the sample period. These trades are denoted as Time Delta and are used in Model (3) in Table 3.

The goal of this exercise is to try to remove one possible explanation for the patterns found above. As noted, these patterns could occur because of trading or quotation activity across venues. If these patterns are due primarily to trading activity taking place very close to the observed off-exchange trade, then removing trades in the sample that occur near other trades should eliminate this pattern. If the pattern remains, then it may suggest the findings are instead due to quotation activity.

Model (3) includes only those trades in the Time Delta sample. As in Model (2), all the coefficients prior to $t = 0$ are positive and statistically significant. This suggests that even for isolated trades, odd-lot liquidity is lower at $t = 0$ than at any earlier offset time examined. Unlike Models (1) and (2), however, the coefficients after $t = 0$ are also all positive and statistically significant, indicating that odd-lot liquidity is also higher at all times examined after the trade occurred. As with Model (1), the level of odd-lot liquidity declines and recovers quickly: There was an average of \$37.52 more odd-lot liquidity on-exchange one millisecond before a trade occurred, and an average of \$32.35 more one millisecond after the trade occurred. The model again has trade-level fixed-effects and generates an R^2 of about 75%.

Model (3) once again indicates that on-exchange odd-lot liquidity is significantly influenced by off-exchange trading. The Time Delta sample suggests that this pattern is unlikely due solely to correlated trading across different exchanges and may also be due to quotation activity. Models (1) - (3) all show that on-exchange odd-lot liquidity is significantly lower at the time of the trade than it was 1 or 5 milliseconds before the trade. This takes as given that off-exchange trades should not interact with exchange liquidity without some intermediary working to do so. Further, Models (1) and (3) show that lost liquidity nearly recovers to the level it was 1 millisecond before $t = 0$ within 1 to 2 milliseconds after $t = 0$.

Models (4) - (6) in Table 3 examine the value of odd-lot liquidity on the same side as the trade for the same three datasets as in Models (1) - (3). For off-exchange buys (sells), this sums together the level of on-exchange odd-lot buys (sells). These trades are not called 'eligible' because if the off-exchange trades were sent to exchanges, they would not be able to execute against these orders. However, it is possible that on-exchange odd-lot liquidity would increase if a dealer who facilitates a buy (sell) off-exchange decides to reverse that trade by submitting an on-exchange buy (sell) order that is posted to the exchange as quote. Given the results in Models (1) - (3), this strategy seems possible.

The coefficients for offsets before $t = 0$ in Models (4) - (6) are largely negative and statistically significant. This is particularly true as you get closer to $t = 0$. For all three models, all coefficients 500 milliseconds before and after the trade occurred are negative and statistically significant. For all models, negative and statistically significant coefficients indicate that same side odd-lot liquidity was larger at the time of the trade ($t = 0$) than at the offset time. All three models are fit with trade-level fixed-effects.

As event time approaches $t = 0$, the coefficients in each column remain roughly flat or decline in magnitude. At 5 milliseconds (1 millisecond) before the trade time, there is approximately \$2,873 (\$2,627) less odd-lot liquidity than at the time of the trade for All Data. For the Fixed NBBO sample, this same value is \$18.10 (\$14.54,) at 5 milliseconds (1 milliseconds) before the trade. Finally, for isolated trades in the Time Delta sample, odd-lot liquidity is \$342 (\$331) lower at 5 milliseconds (1 milliseconds) before the trade. Coefficients after $t = 0$ are generally negative, except in Model (2) where they remain positive for approximately 10 milliseconds before moving negative. The consistency with which on-exchange same-side odd-lot liquidity increases at the time of an off-exchange trade ($t = 0$), but is lower before the trade, may suggest a strategy where dealers operating off-exchange use exchanges to offset their off-exchange dealings. These odd-lots appear to be sent to exchanges and survive for only about 1 millisecond.

The results in Table 3 provide some clarity as to how off-exchange dealers may facilitate trades, and how off-exchange trading interacts with on-exchange liquidity. Dealers could employ several strategies, three of which the examined data could support. First, they could act as an intermediary, buying (selling) on-exchange when they receive an off-exchange buy (sell). This would lead to a decline in on-exchange sells (buys). Second, but related, they may aim to change their inventory position, providing multiple quotes on exchanges. They could then receive an off-exchange order, facilitate it, and cancel their on-exchange dealers. Finally, they could facilitate a trade and then use exchanges to reverse that trade. Since most off-exchange transactions are odd-lot sizes, it is natural to look for odd-lot quotations on-exchange when searching for this type of activity. This would lead to an increase in on-exchange buy (sell) orders when we observe an off-exchange buy (sell).

The findings support all three strategies to some degree. However, the process through which dealers may be using exchanges to reverse transactions they facilitated off-exchange (the last three columns of Table 3) appear stronger, more uniform, and with coefficients that are an order of magnitude larger. This process increases the number of on-exchange odd-lot quotes at the time of a trade (within the same millisecond) around an off-exchange trade. This data cannot say for certain if these trades came before or after the off-exchange trade given the relatively course 1 millisecond precision of the TRF timestamps. However, the most likely strategy appears to be that dealers facilitate trades off-exchange and then submit offsetting on-exchange odd-lot quotations with the hope of receiving high quality trade executions which replenish their inventory.

Finally, it is worth noting that the generally positive coefficients immediately before a trade in Models (1) - (3) and the negative coefficients immediately before a trade in Models (4) - (6) indicate that trading seems to affect both sides of the book. Regardless of whether a trades was a buy or a sell, this table suggests that the trade facilitation process likely impacts both sides of exchanges' odd-lot order books.

The next set of trade-level analyses regress OLPI on dummies for time offsets. This model also includes trade-level fixed-effects.

$$OLPI_{s,i,t} = \alpha_i + \sum_{h=-H}^{H} \beta_h D_{i,t}^h + \epsilon_{s,i,t} \tag{7}$$

This model is used across the three primary datasets examined in this study and the same data sets that are included in Table 3. Since OLPI is dependent upon the side of the trade, there is no differentiation between eligible and same side quotes. For this reason, there are half as many regression outputs.

The results for these analyses are in Table 4, which lists the samples at the top of the table. The offsets are in seconds, with the contemporaneous $t = 0$ offset set as the excluded baseline.

Table 4: Regression of OLPI on Offsets, Percent of Trade Value (bps.)

Offset (sec.)	(1) All Data	(2) Fixed NBBO	(3) Time Delta
-5	0.049	0.011	-0.026
	(46.80)	(8.13)	(-17.47)
-2	0.021	0.007	-0.017
	(25.69)	(8.69)	(-17.97)
-1	-0.001	0.007	-0.017
	(-2.17)	(13.75)	(-18.42)
-0.5	-0.012	0.002	-0.014
	(-26.96)	(12.17)	(-28.52)
-0.4	-0.013	0.001	-0.012
	(-30.55)	(8.50)	(-28.16)
-0.3	-0.013	0.001	-0.011
	(-35.27)	(7.36)	(-29.03)
-0.2	-0.011	0.001	-0.009
	(-31.42)	(6.67)	(-27.38)
-0.1	-0.015	0.001	-0.006
	(-49.10)	(10.82)	(-27.79)
-0.05	-0.015	0.001	-0.005
	(-57.38)	(12.78)	(-26.85)
-0.04	-0.016	0.001	-0.004
	(-59.61)	(13.14)	(-26.58)
-0.03	-0.016	0.001	-0.004
	(-63.10)	(13.95)	(-26.28)
-0.02	-0.016	0.001	-0.004
	(-63.52)	(13.96)	(-11.52)
-0.01	-0.015	0.001	-0.003
	(-72.55)	(14.07)	(-24.93)
-0.005	-0.013	0.000	-0.002
	(-71.24)	(13.32)	(-18.90)
-0.001	-0.005	0.000	-0.000
	(-45.78)	(10.87)	(-0.10)
+0.001	-0.049	-0.000	-0.005
	(-315.83)	(-9.06)	(-60.48)
+0.005	-0.082	-0.001	-0.007
	(-386.90)	(-12.79)	(-66.51)
+0.01	-0.088	-0.001	-0.008
	(-392.88)	(-13.12)	(-65.35)
+0.02	-0.095	-0.001	-0.008
	(-396.72)	(-11.36)	(-62.36)
+0.03	-0.099	-0.001	-0.008
	(-394.99)	(-10.90)	(-59.64)
+0.04	-0.101	-0.001	-0.008
	(-391.93)	(-10.29)	(-54.95)
+0.05	-0.103	-0.001	-0.008
	(-381.61)	(-9.28)	(-52.28)
+0.1	-0.110	-0.000	-0.009
	(-372.38)	(-6.91)	(-44.46)
+0.2	-0.118	-0.000	-0.012
	(-361.19)	(-4.13)	(-40.20)
+0.3	-0.123	-0.000	-0.014
	(-342.00)	(-2.55)	(-38.24)
+0.4	-0.127	-0.000	-0.015
	(-330.46)	(-1.01)	(-36.61)
+0.5	-0.131	0.000	-0.017
	(-326.19)	(2.78)	(-36.45)
+1	-0.144	-0.024	-0.025
	(-289.12)	(-50.93)	(-37.72)
+2	-0.161	-0.067	-0.033
	(-238.50)	(-79.70)	(-34.91)
+5	-0.185	-0.135	-0.044
	(-189.45)	(-101.34)	(-29.38)
R^2	0.80	0.93	0.92
Obs.	315,565,043	105,282,791	81,527,288
Trade F.E.	Yes	Yes	Yes

This table contains the coefficients (t-stats.) for the regression of counterfactual OLPI on time offsets. Data are collected from the LSEG Thesys database covering trades on the TRF on January 16th and 19th, 2024 between 9:35 am and 3:55 pm. Offsets are seconds before (negative) or after (positive) the time that the trade occurred (offset of 0.0).
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

The coefficients are the average percentage of price improvement in basis points which could be generated from executing against the best available odd-lot liquidity on exchanges relative to the OLPI of trade received at $t = 0$. Positive coefficients indicate that the offset time had odd-lot liquidity available which would have improved the prices of the trade for the investor while negative coefficients indicate that the best available odd-lot liquidity would not have improved pricing.

Model (1) contains the results for All Data in the sample. The coefficients are generally negative and statistically significant, indicating that OLPI was on average lower before a trade occurred than it was at the time of the trade. There are exceptions to this pattern, specifically when the offset is greater than one second before the trade. Odd-lot liquidity would have improved the average trade's price by about 0.05 bps had it been executed 5 seconds earlier, relative to that at $t = 0$. As we get closer to $t = 0$, these coefficients become and remain negative. At 5 milliseconds (1 millisecond) before the trade, OLPI was lower by about 0.013 (0.005) basis points. These results, however, are relative to a baseline OLPI that is already negative. Therefore, OLPI is higher at $t = 0$ than it was previously, but it still does not represent an opportunity for investors to outperform by executing against odd-lot liquidity alone. The economic magnitude of these values is also quite small as these represent fractions of a basis point. Therefore, it is likely incorrect to assume these results represent significant value to the average investor or trader.

Instead, these findings can be viewed as a way of measuring the impact of off-exchange trading on the relative value of exchanges' odd-lot order books. Model (1) indicates that within the same millisecond that a trade occurs ($t = 0$), OLPI is higher than it is for most times around the trade. This suggests that the off-exchange trading appears to make on-exchange odd-lot liquidity more valuable for a very brief period, even if the quantity of odd-lot liquidity declines (as shown in Table 3). This suggests that off-exchange trading impacts not only the depth of odd-lot liquidity, but the value of those orders as well.

The Fixed NBBO sample used in Model (2) is smaller but shows a different pattern. In this table, the coefficients are all positive and statistically significant before the trade, but noticeably smaller in magnitude compared to Model (1). This suggests that OLPI and the associated on-exchange liquidity decline in value in the same millisecond as a trade occurs. This model lends support to the idea that, when NBBO remains unchanged, off-exchange dealers may purchase on-exchange odd-lots within the same millisecond that a trade occurred off-exchange. After the trade occurs, the OLPI appears to remain lower than it was at the trade time and remains so for the remainder of the interval. This result is interesting and suggests that the exchange odd-lot order book may not fully replenish itself after an off-exchange trade, suggesting a lasting impact to liquidity. Similar to Model (1), these results are statistically significant, but they are likely not economically significant as the effects are usually less than about 0.01 basis points.

Model (3) repeats this analysis using the Time Delta sample. The results in Model (3) are similar to those in Model (1). For this sample, the pre-trade ($t < 0$) offsets are all negative and statistically significant, except for 1 millisecond prior to the trade. All coefficients after the trade are negative and statistically significant. This suggests that OLPI once again increases at the time of the trade. However, and with the other OLPI results, these values are likely not economically significant to the average investor and are not enough to make OLPI positive. Model (3) shows that the increase in OLPI appears to occur slightly before the trade, but, once again, the overall size of this pattern is small.

These results tell a generally consistent story. On-exchange odd-lot liquidity is likely impacted by off-exchange trading. Regardless of whether a trade was a buy or a sell, the depth of odd-lot bids and asks changes in very close proximity to that trade. However, the average off-exchange trade would likely not have benefited if it were directed to exchanges' odd-lot liquidity instead of being executed off-exchange.

A potential explanation for this is high-frequency dealer price referencing activity. As discussed by Van Kervel & Yueshen (2023), dealers may optimally choose to soften their on-exchange quotes when facilitating trades off-exchange. When dealers price reference to learn more about the value of on-exchange liquidity, they may adjust their own quotations on-exchange, potentially leading to a decline in on-exchange odd-lot liquidity values around a trade. This process could explain the decline in on-exchange odd-lot liquidity if dealers work to do so at high frequency.

The results in Tables 3 and 4 suggest that high-frequency price referencing may occur, wherein dealers facilitating trades off-exchange adjust their liquidity on-exchanges during the trade facilitation process. However, Table 4 suggests that this finding may not significantly impact the value of trade prices received by investors. Further, the results in Table 3 could also be consistent with a story in which dealers capture odd-lot liquidity to manage their inventories.

Overall, the event studies conducted in this section suggest that off-exchange trading has a significant impact on the amount of odd-lot liquidity available on exchanges. Odd-lot liquidity can fluctuate significantly even within the same millisecond that a trade occurred. The patterns reviewed above are consistent with a story in which some off-exchange liquidity providers effectively use exchanges to reverse the transactions they facilitate off-exchange. They can do this by either submitting marketable orders that absorb liquidity on the opposite side of the exchange order book as the incoming off-exchange trade or by posting orders to exchanges in the hope of getting a better price. This pattern suggests that off-exchange trading is actually quite integrated into the national market system. However, off-exchange trading still appears to provide investors with high-quality trade executions, providing prices that are better on average than the prevailing odd-lot liquidity prices.

4.3 Order Cancellations

Evidence provided thus far supports a strong connection between on-exchange quotation activity and off-exchange trading. Quote depths decrease on the side of the book that incoming trades could have been paired with, but increase in value on the opposite side. In this section, a potential mechanism is explored which could explain this pattern.

There are three primary ways odd-lot liquidity can decline in value: (i) a lack of odd-lot quotation activity, (ii) execution of odd-lots against marketable orders, or (iii) cancellation of odd-lot liquidity. It has been shown that the total value of odd-lot liquidity actually increases significantly around a TRF trade, so there does not appear to be a lack of quotation activity. The change in odd-lot liquidity should therefore result from either an increase in correlated on-exchange marketable orders around a TRF trade or from cancellations of odd-lot quotes.

Correlated trading and order cancellations can both affect odd-lot liquidity because they both decrease the number of shares available on-exchange. Decreasing on-exchange odd-lot liquidity influences the relative price improvement that dealers provide to those trading off-exchange. However, correlated trading is more likely to draw scrutiny. If dealers were to trade around TRF trades, they could impact on-exchange prices, which could allow them to transact off-exchange at prices better for them. However, trading before or around a TRF trade is likely too risky of a strategy for large dealers.

Order cancellations can also change the value of odd-lot liquidity on-exchange, but they may be less likely to draw scrutiny. First, they are common practice: There are more order cancellations than trades or unique NBBO quotations. Second, they are also more difficult to examine because they are not recorded in commonly used databases.

In fact, canceling orders is likely an important part of dealers' trade-facilitation process. Consider a dealer that actively provides quotations on-exchange. This dealer may simultaneously facilitate off-exchange transactions, providing execution services directly to clients that do not wish to transact on an exchange. If the dealer provides a quotation on-exchange which has not yet executed and then receives an off-exchange order for the same stock on the opposite side (buy versus sell) and for a similar share size, the dealer could choose to internalize the off-exchange transaction. Canceling the on-exchange trade may be how the dealer transfers the on-exchange liquidity provision to their off-exchange dealings. This could soften on-exchange quotations.

To examine this process, order cancellations are collected for the 100 sample stocks across all trading days analyzed. The raw data is too voluminous to analyze efficiently because it contains approximately 60 cancellations for every one order in the high-frequency sample of trades. For this reason, only the odd-lot order cancellations (cancellations of orders for less than 100 shares) that occur 50 milliseconds before and after a TRF trade occurs are used.

Figure 3 plots the average number of cancelled shares around a TRF trade in ten-millisecond buckets. In this graph, the TRF trade occurs at time $t = 0$ on the x-axis. Negative times are cancellations that occur before the TRF trade while positive times signify cancellations after the TRF trade. Cancellations that occur after the TRF trade should have no impact on the TRF trade's price. However, cancellations that occur before could indicate that dealers are canceling their on-exchange orders prior to facilitating an off-exchange transaction.

Figure 3 suggests that there is a noticeable increase in the number of odd-lot order cancellations on-exchange in the milliseconds before an off-exchange trade occurs. This increase becomes most noticeable about ten milliseconds before a TRF trade. There are also a significant number of order

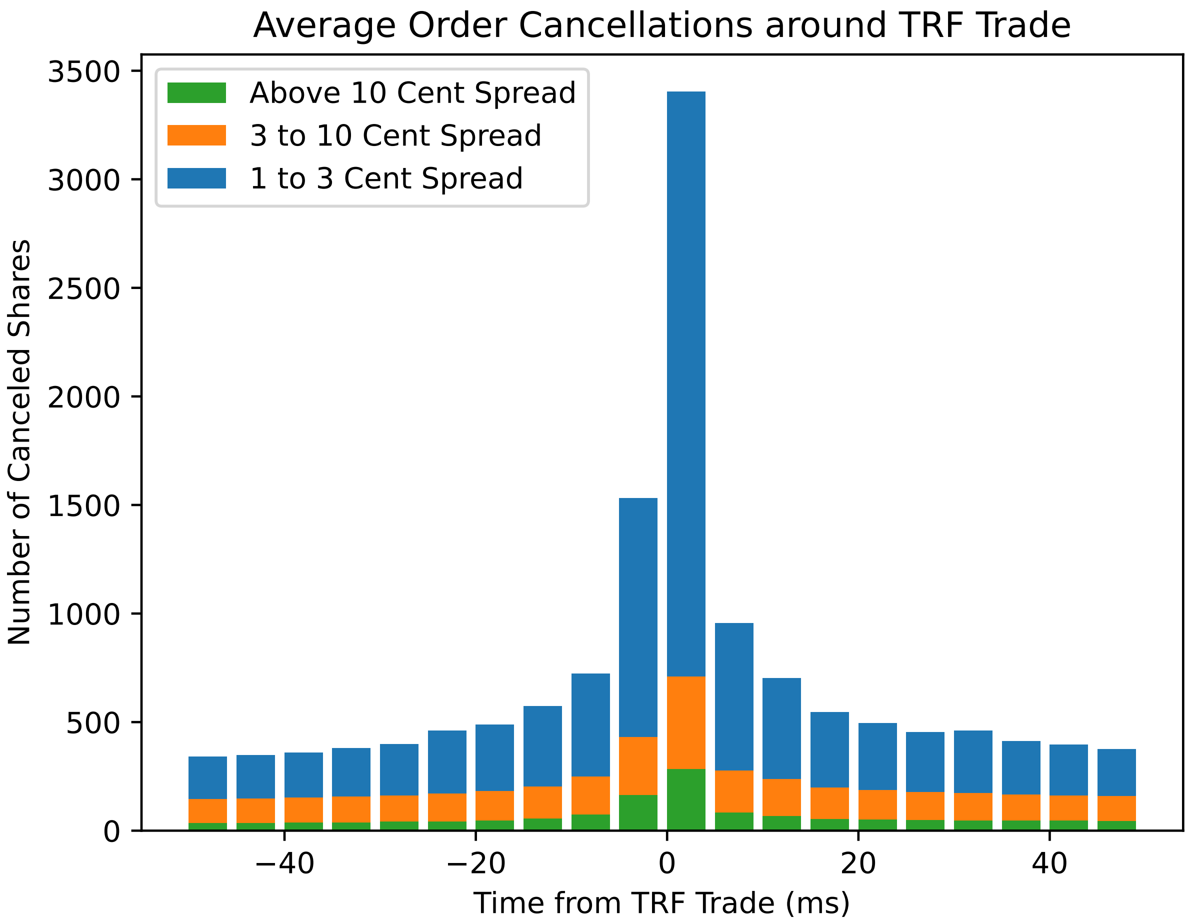

Figure 3: Order cancellations

Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

cancellations which occur after the TRF trade, but those are unlikely to affect the price of the trade. These trades could be, in part, due to a measurement issue. The TRF reports timestamps only to the millisecond while exchanges generally report to the microsecond. If TRF timestamps are truncated and not rounded, as seems to be the likely policy, approximately one millisecond worth of cancellations of the first post-trade bar (the tallest bar) could have actually occurred before the TRF trade. No manual reclassification of these trades is attempted to avoid overestimating the size of order cancellations.

The data is now used to test whether the uptick in order cancellations around a TRF trade is statistically significant. The on-exchange order cancellations data from Figure 3 are used to predict an incoming TRF trade. The number of trades that occur in each ten-millisecond window

Table 5: Purchases and Sales, Cancellations

	All Trades (1)	Purchases (2)	Sales (3)
$BuyCancels_{t-1}$	0.008	0.004	0.004
	(8.599)	(6.193)	(10.560)
$SellCancels_{t-1}$	0.008	0.004	0.003
	(7.795)	(10.385)	(5.095)
R^2	0.008	0.005	0.004
Obs.	456,000,100	456,000,100	456,000,100
Fixed Effects	Symbol	Symbol	Symbol
SE Clusters	Yes	Yes	Yes

This table contains the coefficients (t-stats.) for the regression of the total number of trades which occurred in a ten-millisecond window on the number of buy and sell cancellations measured in number of shares, in the previous ten-millisecond window. Separate intercepts are used for each stock symbol. Data are collected from the LSEG Thesys database covering trades on the TRF on January 16th and 19th, 2024 between 9:35 am and 3:55 pm.
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

is regressed on the number of buy and sell order cancellations over the previous ten-millisecond window.

This analysis serves three purposes. First, it helps answer an important question: Do order cancellations precede trades? This analysis predicts incoming retail orders using the number of order cancellations which occurred shortly before the TRF trade. Second, this setup helps to offset the effects of increasing trading activity. Perhaps order cancellations and retail orders tend to occur together. The lag-lead relationship helps account for this, as do autocorrelation-adjusted standard errors. Finally, this analysis tests if early order cancellations are specific to the type of incoming order. That is, do dealers more frequently cancel their purchase orders before an incoming TRF sell, or vice versa?

Table 5 contains results for all order cancellations and not just those for odd-lot orders. In Model (1), the total number of trades that occurred in any given ten-millisecond window is regressed on the number of buy and sell cancellations that occurred in the previous ten-millisecond window. Stock fixed-effects help control for stock-specific cancellation patterns. Standard errors are clustered and adjusted for autocorrelation.

The results suggest that both buy and sell cancellations have the same effect on the number of trades which occur in any given ten-millisecond window. Specifically, as the number of incoming

Table 6: Purchases and Sales, Odd-Lot Cancellations

	All Trades (1)	Purchases (2)	Sales (3)
$BuyCancels_{t-1}$	0.042	0.023	0.019
	(11.737)	(9.712)	(11.662)
$SellCancels_{t-1}$	0.041	0.018	0.022
	(10.033)	(10.505)	(7.929)
R^2	0.008	0.005	0.004
Obs.	456,000,100	456,000,100	456,000,100
Fixed Effects	Symbol	Symbol	Symbol
SE Clusters	Yes	Yes	Yes

This table contains the coefficients (t-stats.) for the regression of the total number of trades which occurred in a ten-millisecond window on the number of odd-lot buy and sell cancellations in the previous ten-millisecond window. Separate intercepts are used for each stock symbol. Data are collected from the LSEG Thesys database covering trades on the TRF on January 16th and 19th, 2024 between 9:35 am and 3:55 pm.
Source: Muzan TAQ, LSEG Thesys, Author's Analysis.

buy and sell order cancellations increases by 1,000 shares, or 10 round lots, we should expect approximately eight more off-exchange trades to occur within the next window. This suggests that we can predict incoming off-exchange TRF orders using the number of on-exchange order cancellations.

Models (2) and (3) separate TRF purchases and sales. In both, the effect of on-exchange buy and sell cancellations is positive and statistically significant. However, for both, the opposite-side effect has a larger t-statistic, indicating that there is a more reliable effect.

In Model (2), for every 1,000 additional shares of buy or sell orders cancelled, we would expect four more purchases in any given 10-millisecond period. In Model (3), for every additional 1,000 shares cancelled on the buy side, four more purchases are expected on the TRF. For every additional 1,000 shares cancelled on the sell side, three more purchases are expected.

Table 6 contains the same analysis as in Table 5 but uses only the odd-lot order cancellations on-exchange. The results are also everywhere positive and statistically significant, but the coefficients tend to be an order of magnitude larger. For every 1,000 additional shares of on-exchange odd-lot buy order cancellations, 42 additional off-exchange trades are expected while, for sell order cancellations, 41 are expected.

Separating off-exchange buys and sells, the same pattern is repeated. Off-exchange buys and sells are both significant, but the opposite-side transaction has a larger t-statistic. However, as with Model (1), using odd-lot liquidity alone generates significantly larger coefficients. Order cancellations in the ten milliseconds before a trade increase the number of expected TRF trades in the following ten milliseconds by between 18 and 23 trades per 1,000 shares cancelled.

The results in this section highlight an important pattern. Order cancellations appear to rise *before* an off-exchange trade occurs. In a regression setting over ten-millisecond intervals, an increase in order cancellations does precede trades to a statistically significant extent. There is no strong evidence that dealers or market participants knowingly cancel cross-side trades (canceling on-exchange bid when they facilitate on off-exchange sell), but the standard errors associated with these cross-side cancellations are lower.

This pattern is consistent with trade predictability. Price referencing, which, in the prior analyses, occurs when a dealer working off-exchange modifies their on-exchange quotes to justify lower-quality trade executions, may imply modification of on-exchange activities around off-exchange trading. No strong evidence is found that dealers cancel only enough shares to offset a corresponding off-exchange transaction.[10]

5 Commentary and Discussion

This paper contributes to a debate over the accessibility and value of odd-lot liquidity. Odd-lots have been recognized as valuable to investors, but the definition of the NBBO excludes these from dissemination. This means that many investors that do not have expensive proprietary data feeds would be unable to observe the value of odd-lot liquidity before making a trading decision.

This problem is particularly relevant for small retail investors. With small order sizes, retail investors have the highest likelihood of executing at improved prices against on-exchange odd-lot liquidity. And yet, it is also unrealistic to assume that these likely unsophisticated traders would have access to expensive proprietary data. Those traders who could benefit most frequently from odd-lot liquidity are probably those who have the lowest ability to observe it.

[10] For example, if a dealer was facilitating an off-exchange buy order for 100 shares, we might expect them to cancel an on-exchange sell order for 100 shares. However, working within computational limitations, no such pattern is found.

In 2020, the SEC proposed the Market Data Infrastructure Rule (MDIR).[11] This rule would incorporate information on odd-lot liquidity by changing the definition of round-lot. Instead of using a 100-share minimum, round-lots would instead be defined on a graduated scale based on the share price. Higher-priced stocks would have a lower threshold to be considered a round-lot. This should have the effect of including limit orders that are now considered odd-lots into the NBBO quote.

While this change would likely capture a portion of the total odd-lot liquidity in the market, this paper highlights that order cancellations could still cause issues. Results show that in certain cases, odd-lots are not very indicative of actionable liquidity. Instead, they are small quotes that are added and cancelled at high speed, making them almost impossible for most traders or investors to use in their decision process. This would result in NBBO displaying prices that may disappear before they could be used.

6 Conclusion

As alternative, off-exchange trading venues continue to gain market share, it is important to consider the impact they have on exchanges. Since exchanges provide valuable pricing services which off-exchange venues do not, ensuring that off-exchange trading does not impede the ability of exchanges to set prices efficiently should be a primary concern. The results in this paper suggest that exchange liquidity, particularly for odd-lots which generate price improvement, appears significantly influenced by off-exchange trading activity.

Odd-lot liquidity increases in size but decreases in relative value in a small window around off-exchange trades, consistent with a movement away from the midpoint. This transition occurs rapidly, in some cases only a few milliseconds before an off-exchange trade occurs, and on average returns to normal levels within half a second. The decrease in odd-lot liquidity is linked to an increase in order cancellations. By design, dealers frequently post quotations to market centers and likely cancel those orders as the spread moves and new information becomes available. Order cancellations appear to arrive just before a trade, but this may be due to highly correlated trading

[11] `https://www.sec.gov/rules-regulations/2020/12/market-data-infrastructure`

on the TRF. These spillovers appear to impact exchange liquidity. This pattern is found to be consistent enough to predict the arrival of a TRF trade by looking at the increase in on-exchange order cancellations about ten milliseconds before the off-exchange trade occurs.

Overall, this paper finds that off-exchange trading likely impacts prices and liquidity on exchanges. While these alternative venues are marketed, reasonably so, as low-cost, high-speed destinations for equity orders, they certainly do not operate in isolation. Instead, the trading activity occurring off-exchange appears to affect the depth and quality of liquidity on-exchanges and does so at very high speed.

References

Baldauf, Markus, Mollner, Joshua, & Yueshen, Bart Zhou. 2024. Siphoned apart: A portfolio perspective on order flow segmentation. *Journal of Financial Economics*, **154**, 103807.

Barber, Brad M, Odean, Terrance, & Zhu, Ning. 2008. Do retail trades move markets? *The Review of Financial Studies*, **22**(1), 151–186.

Bartlett, Robert P, McCrary, Justin, & O'Hara, Maureen. 2022. The market inside the market: Odd-lot quotes. *Available at SSRN 4027099*.

Bartlett III, Robert P. 2021. Modernizing odd lot trading. *Colum. Bus. L. Rev.*, 520.

Battalio, Robert, & Holden, Craig W. 2001. A simple model of payment for order flow, internalization, and total trading cost. *Journal of Financial Markets*, **4**(1), 33–71.

Battalio, Robert, Corwin, Shane A, & Jennings, Robert. 2016. Can brokers have it all? On the relation between make-take fees and limit order execution quality. *The Journal of Finance*, **71**(5), 2193–2238.

Battalio, Robert H. 1997. Third market broker-dealers: Cost competitors or cream skimmers? *The Journal of Finance*, **52**(1), 341–352.

Bessembinder, Hendrik, & Kaufman, Herbert M. 1997. A comparison of trade execution costs for NYSE and NASDAQ-listed stocks. *Journal of Financial and Quantitative Analysis*, **32**(3), 287–310.

Boehmer, Ekkehart, Jones, Charles M, Zhang, Xiaoyan, & Zhang, Xinran. 2021. Tracking retail investor activity. *The Journal of Finance*, **76**(5), 2249–2305.

Dorn, Daniel, Huberman, Gur, & Sengmueller, Paul. 2008. Correlated trading and returns. *The Journal of Finance*, **63**(2), 885–920.

Easley, David, Kiefer, Nicholas M, O'hara, Maureen, & Paperman, Joseph B. 1996. Liquidity, information, and infrequently traded stocks. *The Journal of Finance*, **51**(4), 1405–1436.

Ernst, Thomas, Spatt, Chester S, & Sun, Jian. 2022. Would Order-By-Order Auctions Be Competitive? *Available at SSRN 4300505*.

Harris, Lawrence. 1989. A day-end transaction price anomaly. *Journal of Financial and Quantitative analysis*, **24**(1), 29–45.

Hu, Edwin, & Murphy, Dermot. 2024. Competition for retail order flow and market quality. *Available at SSRN 4070056*.

Kumar, Alok, & Lee, Charles MC. 2006. Retail investor sentiment and return comovements. *The Journal of Finance*, **61**(5), 2451–2486.

Lee, Charles MC, & Ready, Mark J. 1991. Inferring trade direction from intraday data. *The Journal of Finance*, **46**(2), 733–746.

Mahoney, PG, & Rauterberg, GV. 2018. Chapter 5. The Regulation of Trading Markets. *Securities Market Issues for the 21st Century. Prepared in connection with The New Special Study a project of Columbia Law School & Columbia Business Schools Program in the Law and Economics of Capital Markets. Independently published*, 222–278.

O'Hara, Maureen, Yao, Chen, & Ye, Mao. 2014. What's not there: Odd lots and market data. *The Journal of Finance*, **69**(5), 2199–2236.

Schwarz, Christopher, Barber, Brad M., Huang, Xing, Jorion, Philippe, & Odean, Terrance. 2023. The 'Actual Retail Price' of Equity Trades. *SSRN*, **Working Paper**.

Van Kervel, Vincent, & Yueshen, Bart Zhou. 2023. Anticompetitive Price Referencing. *Available at SSRN 4545730*.