Deviations from Time Priority on the NYSE

Robert Battalio
Professor of Finance
Mendoza College of Business
University of Notre Dame

Robert Jennings
Professor of Finance (Emeritus)
Kelley School of Business
Indiana University

Bill McDonald
Professor of Finance
Mendoza College of Business
University of Notre Dame

First Draft: February 2018

Current Draft: March 8, 2019

Abstract: Most financial security trading venues prioritize competing liquidity providers with price-time priority; entities displaying the best-priced limit order earliest trade with the next opposite-sided liquidity-demanding order. The New York Stock Exchange (NYSE) rule putting traders "on parity" is a vestige of the traditional floor-based trading model and likely favors floor-based traders. Parity requires that orders from floor brokers, the designated market maker, and the top of the electronic limit order book trade together. Thus, floor trading interests can trade ahead of equally priced, previously arriving orders in the limit order book. The NYSE posits that floor traders provide valuable services to the investing public and one approach to remunerating them for these services is deviating from strict time priority. The cost of the potential benefits of the NYSE's business model is opaque. Our research is an initial attempt to quantify this cost.

We thank workshop participants at Indiana University and Notre Dame University and Shane Corwin, Adam Nunes at Hudson Trading and Hank Bessembinder for comments.

Even in today's electronic environment, the New York Stock Exchange (NYSE) maintains remnants of their traditional floor-based trading. Specifically, while other equity trading venues have abandoned trading floors, NYSE rules provide for two types of floor-based traders – Designated Market Makers (specialists in prior lexicon) and floor brokers. The NYSE argues that these floor participants provide benefits to investors and grants specialized trading arrangements to compensate them for this public service. Of interest to this paper, floor participants potentially can trade ahead of orders on the public limit order book arriving earlier. That is, strict time priority is not enforced on the NYSE. The NYSE's compensation mechanism represents a transfer payment between various categories of traders and is not transparent. We estimate the size of the transfer payment associated with this particular aspect of the NYSE's business model in our paper.

The NYSE's competitors use price-time priority to arrange trades.[1] That is, liquidity-supplying traders willing to display the most favorable price to incoming liquidity-demanding orders earliest get the opportunity to trade with the next arriving opposite-sided liquidity-demanding order. For liquidity-supplying buys (sells), the limit order offering a liquidity-demanding seller (buyer) the highest (lowest) price among all liquidity suppliers is first in line. To break ties among interests on the limit order book with the same price, most venues use time-of-arrival as a secondary precedence rule; orders displayed sooner have priority over equally-priced orders arriving later. Price-time priority is hypothesized to give traders an incentive to place aggressively priced limit orders, which produce tight quoted spreads that attract liquidity-demanding trades and offer traders lower transaction costs (see, e.g., Harris (2003) and Hasbrouck (2007)). Although Securities and Exchange Commission Regulation National Market System (Reg. NMS) Rule 611 enforces price priority in U.S. equities markets, liquidity-

[1] There are minor exceptions to this statement. For example, Nasdaq PSX trades on a price-quoted size priority.

demanding traders are not required to respect time priority across execution venues and common order routing practices routinely violate across-market time priority.

The New York Stock Exchange (NYSE) uses price-parity to arrange trades amongst market participants. Although public orders on the electronic order book use price-time priority, incoming marketable orders are allocated between the public book, the designated market maker (DMM) and floor brokers committed to trading at that price.[2] This mimics the traditional floor model where trading was apportioned between the public order book and the "crowd" of floor-based traders. Under most trading conditions, parity benefits floor participants (and their customers) at the expense of non-floor traders with orders on the NYSE's public order book. When public limit orders arrive, they are placed in time priority on the (potentially deep) book. When floor limit orders arrive, they are on parity with the top of the public limit order book. Thus, DMMs and floor brokers do not have to wait in the public order queue to interact with marketable orders and, potentially, can strategically use parity to mitigate the adverse selection problem associated with displaying trading interests in long queues at the expense of non-floor investors.

NYSE literature states that there are four benefits associated with using price-parity to arrange trades. First, the NYSE asserts that the parity model levels the playing field by allowing institutional and retail traders to share executions at the same price as the one quoted by faster market participants. Second, it argues that parity delivers higher fill rates and lower transactions costs for institutional and retail investors. Third, the NYSE notes that parity encourages broad participation from all customer segments. The fourth claim is that parity results in deep liquidity

[2] See NYSE Rule 72 or https://www.nyse.com/publicdocs/nyse/markets/nyse/Parity_and_Priority_Fact_Sheet.pdf. Priority gives orders that set the best NYSE quote a guaranteed 15% share (minimum one round lot) of each execution at that price regardless of parity. Priority rewards traders bettering the NYSE quoted price and reduces parity's effect.

and superior market quality. In this paper, we use the recently made publicly available TAQ NYSE XDP Integrated Feed to empirically explore the costs associated with these potential advantages.[3]

There is substantial practical evidence that the NYSE's floor-based trading model bestows benefits to and imposes costs on different sets of investors. NYSE floor brokers advertise their ability to trade at parity to attract order flow, which suggests that trading at parity has value. For example, Meridian Equity Partners notes that "clients who trade through Meridian's NYSE floor brokers obtain parity executions" and that "client orders bypass the limit book queue, trade equally with/alongside other market participants thus allowing for superior liquidity and fill rates."[4] That orders resting on the public limit order book are disadvantaged by parity is noted by a Canadian investment firm who states that when a DMM or floor broker steps ahead of an investor's order, "the investor's only solutions are to not trade or to remove the quote from the limit order book and become active." This firm notes that both scenarios result in substantially higher transaction costs for orders resting in the public order book.[5] Evidence presented in Battalio et al. (2016) suggests that few passive retail orders are routed to the NYSE. Thus, it is likely that the passive orders resting on the NYSE's public order book are placed there by algorithms programed to account for the effects of parity when making order routing decisions. Whether parity allocations impose large costs on actively-managed passive orders or whether sophisticated order routers make such costs minimal is an empirical issue we address in this paper.

[3] See https://www.nyse.com/market-data/real-time/integrated-feed. The data are first available in late 2015.
[4] See http://www.meptraders.com/ElectronicExecution.html.
[5] See http://www.osc.gov.on.ca/documents/en/Marketplaces/com_20140911_drakej.pdf.

Unlike DMMs, floor brokers do not trade proprietarily nor do they incur additional obligations to trade at parity.[6] Responding to a recently approved NYSE proposal to use the price-parity model to trade Nasdaq-listed securities, the Chicago Board Options Exchange notes that the "preferential treatment provided to Floor Brokers in a fully electronic trading environment … does, in fact, unfairly discriminate against market participants that do not submit orders through a Floor Broker and provides Floor Brokers with an unnecessary competitive advantage".[7] Given that each floor broker trades at parity in a given stock, their presence can amplify the adverse selection costs imposed on investors who place limit orders in actively-traded stocks with thick limit order books.

We document parity executions on the NYSE and assess the economic implications of parity. Similar to Ready's (1999) analysis of specialists, our thesis is that floor traders take advantage of parity more frequently when being at the front of the queue is valuable and competing mechanisms for obtaining priority are less useful. We anticipate that parity's use is associated with order and stock characteristics and market conditions. We find that traders successfully use parity more frequently for stocks with lower prices and less price volatility. We also document that parity's prevalence increases when the minimum pricing increment of one penny is binding (making obtaining priority via price difficult), when the NYSE has more relevant-sided quoted depth at the best quoted price, when competing venues designed to gain priority are crowded, when a purchase (sale) occurs at a price that is higher (lower) than the prior NYSE trade price, and toward the end of the trading day. These attributes imply that being at the

[6] As floor brokers trade only for clients, their parity allocations represent a benefit to investors (paid in the form of a transfer payment from another market participant). This benefit is available to any investor choosing to use a floor broker. Why investors choose not to use floor brokers to obtain this advantage is a topic outside the scope of this paper. Clearly, there must be some costs of using a floor broker that offset the advantage. Although we understand that the direct commissions charged by floor brokers are not large, it seems that the floor broker trading system is less nimble than some of the fastest traders prefer. Recent research by FINRA (Anand et. al. 2019) suggests that there might be agency costs at work in order routing decisions in addition to direct costs.

[7] See https://www.cboe.com/publish/ComLet/20180201.pdf.

back of the book significantly decreases the chance that a limit order fills or that it fills only when a large (potentially informed) opposite-sided order arrives and/or multiple opposite-sided orders arrive.

We also estimate the gross profit generated by parity trading during May 2017. To the extent that floor participants strategically employ parity, non-floor orders in the limit order book are disadvantaged. Floor-based trades compete with the book when it is profitable to do so but let potentially informed liquidity-demanding orders interact with the public book. To estimate the value of this dynamic with our data, we analyze the prices at which orders bypassed by parity trades fill compared to the price that they would have received if time priority was enforced on the NYSE and estimate the cost associated with bypassed orders that are deleted after being bypassed. We find that orders bypassed by parity trades in NYSE-listed securities fill at worse prices, on average, than they would have received with strict time priority, representing a $56,500 aggregate cost during our sample month. For orders deleted after a parity order steps in front of them, we attempt to bracket our implicit cost metrics. To measure the opportunity costs associated with unfilled bypassed orders, we compare the price at which the order would have traded without parity executions to the end-of-day quoted spread mid-point. That produces an opportunity cost of nearly $710,000. Assuming that *all* traders who delete orders bypassed by parity orders respond by submitting an aggressive order, we estimate a sample-month cost of $3.7 million. Although order routers might recognize the effects of parity trading on the NYSE, this research represents the first explicit measurement of parity's cost to the authors' knowledge.

After illustrating how parity works in the next section, we use subsequent sections to survey relevant extant literature, introduce our data, and explain our methodology. In Section VI, we provide an empirical analysis of the use of parity on the NYSE. Section VII discusses our

estimates of the costs of parity trading to public orders on the electronic book. Section VII concludes.

II. An illustration of parity.

To illustrate more completely how parity influences NYSE trade executions, consider the hypothetical example summarized in Table I. Panel A establishes the top of the buy (bid) side of the NYSE public limit order book at the National Best Bid (NBB) price for a particular point in time. All six orders share price priority. The orders are arranged based on time priority – order 1 was the first to arrive and order 6 the most recent. By construction, the first four orders are public orders and the last two orders are floor orders.

Assume that a 400 share market sell order arrives. With strict time priority, orders 1, 2, and 3 would provide liquidity to the market sell order. Orders 4, 5, and 6 would not trade. With parity, public order 1, floor broker order 5, and Designated Market Maker (DMM) order 6 fill completely and 100 shares of public order 2 executes (leaving 100 shares).

Panel B outlines the sequence of pairings generated by the arrival of the 400 share marketable order under parity. Although these are not trades in the formal reporting sense, they illustrate how parity allocations occur. In the first transaction, 300 shares trade from public order 1 and floor orders 5 and 6 with each liquidity-supplying order receiving 100 shares. That is, parity orders 5 and 6 fill share-for-share (in round lots) with the top of the limit order book. Orders 1, 5, and 6 are completely filled after the first transaction. In the second transaction, public order 2 gets 100 shares leaving 100 shares. After the market order is fully executed, 100 of the original 200 shares in public order 2 and all shares of public orders 3 and 4 remain unexecuted. Thus, public order 2 is 100 shares short of the fill it would have received without

parity and public orders 3 and 4 do not trade and are 200 shares behind where they would have been without parity orders.

In assessing the cost of the parity rule (conversely the benefit to trading at parity), there are several considerations. Order 1 suffers no injury in our example. Order 2 leaves 100 shares that would have executed without parity but remain unfilled with parity. Order 3 is 100 shares short of the fill that it would have received with strict time priority. If orders 2 and 3 eventually fill, we can compute the difference in fill prices received relative to the price that they would have received under strict time priority. Assessing the monetary damage to these orders when they do not fill relies on assumptions about subsequent trader action, which we discuss later.

III. Related literature.

Theoretical work on time priority typically compares the effects of different secondary priority rules (price priority being the primary priority rule). Pirrong (1996) speculates that strict enforcement of time priority might hurt floor traders in futures markets who traditionally trade with the top of the public order book via pro rata allocation rules (all at the same price share in each trade) and reduce their participation. Cordella and Foucault (1999) analytically demonstrate that shifting from strict time priority to a pro rata allocation between all dealers quoting the best price reduces the incentive to compete and widens the bid-ask spread. Panchapagesan (1997) theoretically compares time priority to a rule that allocates trades proportionally to the size quoted. He finds that time priority encourages price competition, narrowing the bid-ask spread.

Because time priority is so common, there are few empirical studies of its value. Panchapagesan (1997) compares competing secondary priority rules on the Toronto Stock Exchange, one using time and another using size. He finds that trading using time priority encourages price competition and results in tighter quoted spreads and more quoted depth. Some

empirical work on the value of time priority examines changes in market structure that alter the strictness with which time priority is enforced. Frino, Hill, and Jarnecic (2002) find no change in bid-ask spreads or price volatility in futures markets when an electronic market using time priority replaces an open outcry system using pro rata allocation. Gwilyn and Alibo (2003) document less price clustering (more price competition) when a floor based market converts to an electronic market, thereby increasing the importance of time priority. Ikenberry and Weston (2008) find more price clustering (less price competition) after equity markets begin trading in decimal increments making it easier to defeat time priority as noted in Harris (1998). Unlike these studies, which examine changes to/from parity in an individual market, we examine how parity is used in a fragmented market where passive investors can avoid being stepped ahead of by rationally choosing between competing venues that have different rules.

A few papers examine the value of stepping ahead of public limit orders in a pre-Reg. NMS (non-electronic) trading environment. Ready (1999) examines the specialist's ability to step in front of the book to execute an order at a better price at a time when the minimum pricing increment was \$0.125.[8] He documents that trades where the specialist immediately steps in front of the limit order book via price improvement were, on average, \$0.059 more valuable than trades where the specialist allows the execution against the limit order book at the prevailing quote. In a related paper, Harris and Panchapagesan (2005) find that NYSE specialists are more likely to improve prices when they can trade in front of a limit order book that has a long queue. Unlike NYSE specialists who could decide whether to trade with a marketable order after it arrived on the floor of the NYSE, in the post-Reg. NMS trading environment DMMs and floor brokers on the NYSE must commit to trade at parity *before* marketable orders arrive. Finally,

[8] At the time, the specialist also had the option of "stopping" an order. The stop guaranteed the marketable order the current best price on the book but gave the specialist the option of improving the price. This process allowed the specialist to observe future order flow to determine if stepping in front seemed warranted.

Moallemi and Yuan (2017) use order data from 2013 (post-Reg. NMS) to estimate the value of being at the front of the limit order book. They do this by comparing the profitability of orders at the front of the queue to the profitability of a hypothetical zero-sized order at the end of the queue. Across the nine sample stocks, their estimated value of priority varies from $0.0005 to $0.0017 per share. We examine actual orders placed with the intention of trading at parity with the NYSE's public order book.

Although dated from periods of time before markets became fully electronically linked post-Reg NMS, extant work suggests that floor brokers provide a valuable service to investors. Werner and Sofianos (2000) find that floor brokers provide two-sided liquidity in many instances where the specialist's participation is limited. Battalio et al (2007) demonstrate that repeated interactions between the same floor brokers and specialists produce lower transactions costs.

Bessembinder et al. (2011) theoretically and Panayides (2007), Clark-Joseph, Ye, and Zi (2017), and Bessembinder et al. (2017) empirically provide support for DMMs' value. Anand and Venkataraman (2016) find that the per-share trading revenue generated by DMMs on the Toronto Stock Exchange (TSX) is higher when they use the 'autoparticipate' feature to trade at parity with the order book.[9] On the TSX, DMMs are the only market participants that can trade at parity and they are allocated approximately 40% of incoming marketable orders when trading at parity. To the extent that trading is a zero-sum game, the fact that DMMs on the TSX earn higher profits when trading at parity suggests that orders resting in the order book pays for their services. Given that all NYSE floor participants (not just DMMs) receive parity allocations, the potential cost imposed by parity trading on the NYSE might be greater than on the TSX.

IV. Data.

[9] The published paper does not report how much higher the profits are. See page 347.

The publicly-available NYSE TAQ XDP Integrated Trades and BBO Feed contains a time-stamped (to the nanosecond) record of all order-related messages transmitted to the NYSE. These data provide details of all open orders and order changes (including modifications, replacements, and deletions), trades, and security status and imbalance messages. For the purposes of our study, we are interested in five message types: add order, modify order, cancel-replace order (hereafter replace), delete order, and order execution. An add order message is published when a new visible order is received and added to the book. A modify order message is published when the price and/or volume of an order is changed due to an event other than a replacement of the original order or full or partial execution of the original order. A replace order message appears when a cancel-replace order is received and completed. A delete order message is generated when an order is taken off of the order book for any reason other than a full execution. Finally, an order execution message is published when the order is fully or partially executed.

An order's history begins with an add order message and can contain multiple modify, replace, and/or execution messages. The order's history ends with a delete message, a complete fill (which can be the accumulation of multiple partial order execution messages), or the end of the trading day. At any time in the trading day, we can determine the price and remaining size of any open order. This might require linking numerous messages modifying, replacing, or partially executing an order. Thus, we can recreate the limit order book – not only knowing the aggregate depth at each price point but also the order in which all orders at a price appeared on the book. With this information, we can determine when orders execute out of time priority (i.e., receive parity allocations). We provide a detailed description of these messages and associated descriptive statistics for the sample month of May 2017 in the Data Appendix.

V. Empirical approach to assessing the effect of trading at parity.

A. Book building.

To determine if some orders appear to execute out of time-priority sequence, we build the book throughout the trading day. At the beginning of the day, we place add order messages on the appropriate side of book at the indicated price level in chronological order. We remove deleted orders from the book at deletion time. Orders executing are decremented by the size of the fill at the time of the trade. A cancel-replace message is treated as a delete message and an add message – the remaining shares of the original order are eliminated from the book and the new order is added to the book at the indicated time. Modifications include changes to price and/or size. We make the indicated changes to the existing order. Order modification messages contain a binary variable that indicates whether the order retains its position in the book. With this approach, we know the exact composition of the book with price and time priority throughout the trading day.

B. Inferring parity trades.

We define a trade as all limit order execution messages with the same nanosecond time stamp. (As noted earlier, a single marketable order can interact with multiple limit orders.) Given the detailed book build outlined above, we examine each trade to determine whether any of the orders in the trade do not follow time priority. We take the number of shares filled from the limit order book as the size of the opposite-sided marketable order.[10] The data provide the shares executed from each individual order in the trade and we know where each order stands in the order book. If the NYSE observed strict time priority across all market participants, then the limit orders executing would fill in the order that they arrived on the book and orders with higher

[10] Sometimes, trades involve an order type allowing a trade at a price within a designated distance from the limit price. Also, there are instances with limit orders on both sides of the trade. We ignore these trades in our analysis.

time priority would fill in their entirety before orders with lower time priority receive executions. Frequently, however, this is not what we observe. Specifically, we find: 1) incompletely filled orders with higher time priority in juxtaposition with lower-priority orders receiving fills, and/or 2) skipped higher-priority orders at the same time orders with lower time priority receive executions.[11]

Two examples from December 7, 2015 posted on the NYSE's website illustrate what we conclude are cases of parity allocations. In the first example, the eight orders in the Ford limit order book with the highest time priority receive executions totaling 18,300 shares at the same time. However, the sixth order in the book leaves shares not filled despite the fact that orders seven and eight receive fills. We infer that orders seven and eight are parity allocations. This violation of strict time priority potentially disadvantages the unfilled shares in order six.

In the second example, 800 shares of Ford stock fill simultaneously. The first order on the book fills 400 of 1,500 shares. The third order on the book fills completely with 300 shares and the fifth order on the book gets all 100 shares wanted.[12] Apparently, orders 3 and 5 trade on parity pro rata with order 1, the order at the front of the queue on the NYSE's limit order book. Here, the first book order is 400 shares behind where it would be with strict time priority. In addition, the second and fourth orders on the book, skipped in their entirety, are 400 shares further in the queue than they would have been under strict time priority. Whether these orders are actually disadvantaged depends on how market conditions differ between when they would

[11] A trade can contain only order(s) out of time priority. This occurs when there are insufficient shares in the liquidity-demanding order of the previous trade to satisfy all parity allocations. For a simple example, consider a situation in which the top of the book (and first in time priority) has a size of 100 shares and there is one DMM order for 100 shares. If a liquidity-demanding order is for 100 shares, the top of the book order fills. If the next liquidity-demanding order also is for 100 shares, the DMM order is "due" and receives a fill even if not first in time priority.

[12] For illustration, think of this 800-share execution as three trades. Trade 1 gives 100 shares each to orders 1, 3, and 5, fully executing order 5. Trade 2 is 200 shares each to orders 1 and 3, fully executing order 3. Finally, order 1 receives the final 100 shares of the trade.

have traded with strict time priority and when they did trade, as well as assumptions about what the trader submitting these potentially disadvantaged orders did if they did not fill.

Our data do not allow us to distinguish who receives parity allocations. Consider a more complex example where correspondence with the NYSE provides additional details. On May 17, 2017 at 14:22:46.667636 an in-sample trade in Ford common stock occurred for 213,206 shares at $10.75. At the time of the trade, the NYSE quote was $10.75 bid for 391,467 shares and 43,060 shares offered at $10.76. The 391,467 shares at the best bid price represented 447 orders.

The execution at least partially filled the first 276 orders in time priority. Specifically, these 276 in-sequence orders received 152,706 shares of the 157,112 shares for which they bid. In addition, there were 12 orders filled out of time priority that we note as parity allocations. The sizes of the 12 orders filled out of sequence, arranged by order arrival time, were four orders of 100 shares each, one order for 1,000 shares, three orders for 100 shares each, one order for 4,900 shares, one order for 21,800 shares, one order for 1,600 shares, and one order for 30,500 shares. Thus, the aggregate amount of parity allocations represented by the 12 out-of-sequence orders was 60,500 shares.[13] There were 101 orders on the book in time sequence bypassed by the 60,500 shares in the 12 out of sequence orders. The first bypassed order had a time stamp of 11:10:46.351322 and the last bypassed order had a time stamp of 13:52:22.139428. The time stamp of the first parity order was 14:16:13.699699 and the time stamp of the last parity order was 14:22:26.686447. Thus, the last (and largest) parity order arrived at the NYSE over three hours after the first bypassed order arrived and roughly 20 seconds before the trade.

VI. Stock characteristics of and market conditions in which parity executions occur.

[13] Conversations with NYSE officials indicate that there were two DMM orders for a total of 700 shares that executed as part of the 276 orders filled in time priority sequence. Our methodology does not allow us to identify these two trades as parity allocations. The NYSE indicates that the 12 orders receiving out of sequence parity allocations consisted of six floor broker orders of 100 shares each and 6 DMM orders totaling 59,900 shares.

As noted above, we define a trade as all execution messages with the same time stamp. Thus, a trade represents the size of the marketable order and, by definition, can involve multiple limit orders. We restrict our analysis to trades executed between 9:45:00 and 3:45:00. Starting at 9:45 abstracts away from complications associated with the opening trade.[14] Ending at 15:45 excludes the Closing Imbalance Period where traders face restrictions on the types of orders that can be submitted.[15] We also eliminate trades with trade prices less than or equal to $1 or trades with trade prices greater than or equal to $1000. Stocks priced below $1 trade on a finer price grid than stocks above that level and stocks with prices exceeding $1000 skew statistical analyses. Table II provides some descriptive statistics for these trades.

In May 2017, our sample of NYSE-listed stocks has a mean trade price of $51.89. The average stock has 43,116 trades exchanging 11,488,297 shares in the month. The range of sample stock characteristics suggests considerable cross-sectional variation. The average trade-weighted relevant-sided quoted depth is 978 shares and approximately two-thirds of the trades occur when the NYSE's quoted spread is $0.01 (e.g., in market conditions where passive investors cannot use price to gain priority on the NYSE). The average marketable order is for 251 shares and interacts with an average of 1.77 passive orders. Approximately 15% of the trades have a parity allocation and passive orders trading at parity represent nearly 10% of the shares traded.[16] On average, there are .18 orders bypassed by an order trading out of strict time priority in each sample trade.

[14] We also eliminate a small number of trades occurring after 9:45 that appear to be the opening trade of the day.
[15] We note that another privilege that floor traders receive is the ability to enter orders intended for the closing trade after the 15:45 cut-off imposed by non-floor traders. See Bacidore et al (2013) for a more complete discussion.
[16] NYSE officials indicate that floor brokers accounted for 5.3% of liquidity providing volume and DMMs provided 16.8% of liquidity providing volume during May 2017. As our methodology cannot detect parity allocations that would execute with strict time priority and not all floor broker and DMM trading is from parity allocations, our estimate should be less than the total floor participation.

Panel B of Table II provides separate descriptive statistics for trades with and without a parity allocation. We expect that parity executions are more likely to occur in actively-traded, low priced stocks where the width of the NBBO is equal to the minimum quoting increment of a penny and limit order queues are long. Consistent with these expectations, on average, trades with parity allocations occur in lower-priced stocks with somewhat higher trading activity. Additionally, the average relevant-side displayed depth is nearly seven times larger for trades with parity allocations and trades with parity allocations occur when the NYSE quoted spread width equals a penny 87% of the time. Finally, conditional on there being a parity allocation, about two-thirds of the marketable order is allocated to parity orders. Not surprisingly, trades with parity allocations tend to be larger both in terms of number of shares traded and in orders receiving fills.

We next investigate whether we can explain the market conditions associated with parity executions in a multivariate setting. We use ordinary least squares (OLS) where the dependent variable is the fraction of shares in a trade allocated to parity and a logistic regression where the dependent variable is binary (one if there is at least one parity allocation in the trade, else zero).

Explanatory variables include liquidity-demanding order attributes, market conditions, and stock characteristics. We control for two attributes of the marketable order: the order's side and the time that the order arrives during the trading day. Of particular interest is the order's timing. We hypothesize that getting to the front of the queue is more important as the trading day nears its conclusion. So, we include the time that the trade occurs (the number of seconds past midnight) and a binary variable that indicates whether the trade occurs between 15:30:00 and 15:45:00.

We measure market conditions with the level of the market's absolute return for the day, a binary variable indicating whether the NYSE quoted spread width is equal to a penny, and the relevant-sided quoted depth on the NYSE. As earlier noted, when the NYSE's spread equals the tick size, parity becomes the only mechanism by which passive investors can gain priority on the NYSE. Similarly, when limit order queues are long, parity allows passive orders on the NYSE to gain priority without improving the limit price by $0.01 per share. Absolute market return controls for the market volatility on the day. As the order book is likely to get deeper the further away limit prices are from the NBBO, we expect more parity trading at the NBB (NBO) after the NYSE's bid price falls (offer price rises). For this reason, we also include a binary variable denoting whether the quoted price associated with the current marketable sell (buy) order is lower (higher) than the quoted price associated with the prior sell (buy) order executed on the NYSE.

The stock attributes we consider are stock price level, stock price standard deviation (volatility), trading intensity (shares traded), and the fraction of trades during the sample month occurring on exchanges that offer liquidity demanders low trading costs (i.e., the so called "inverted" exchanges). We expect time priority to be less important in higher priced stocks as these stocks are likely to have larger price gaps in the NYSE's limit order book. Because liquidity must be exhausted at a price level for prices to move to the next level of the limit order book, we expect parity use to be inversely related to the level of stock price volatility. All else equal, we expect fill rates are positively correlated to trading intensity. Thus, we expect parity usage to be inversely related to trading intensity as measured by the natural log of the total number of shares traded. Finally, we expect parity to be used more frequently if competing approaches to gaining time priority are more crowded. One alternative is to post liquidity-

supplying orders on a venue that has lower access fees. If tied in price priority, then rational liquidity-demanding traders might use these low-fee venues. As a result, we expect the intensity of parity trading on the NYSE to be positively correlated with the intensity of trading on the three low-fee exchanges in May 2017.

In Table III, we report the coefficient estimates from an OLS regression using the percent of shares in a trade that are allocated to parity order(s) as the dependent variable and a logistic regression with a binary variable equal to one if at least one order in the trade is a parity order. The unit of observation is a trade. For a trade in an NYSE-listed stock to be included in the analysis, we require that there be at least 10 trades in the stock in May 2017. This restrictions reduces the sample of NYSE trades to 38,292,542. Given the large number of observations, each of the coefficient estimates is statistically significant at traditional levels.

Our findings regarding market conditions are consistent with our hypotheses; parity allocations are more likely to occur in a given trade if (1) the trade is in a low-priced stock, (2) the market is less volatile on the day and/or the stock itself has low volatility, (3) the bid-ask spread is at the minimum variation, (4) the quoted depth on the relevant side is large, and (5) the trade occurs at a price that is less advantageous for the liquidity demander than the price of the prior trade. In these circumstances, without parity, traders are at the back of a long queue and/or cannot use price priority to gain an advantage. In addition, trades are more likely involve parity when the alternative of going to low-fee venues to pay for time priority is crowded. Trades occurring later in the day – whether measured continuously or discretely – are more likely to include parity, seemingly as traders become more concerned about getting the trade done before day's end. Finally, counter to our intuition, we detect parity more frequently in stocks with more trading during the month.

VII. The value of trading at parity on the NYSE.

Parity's value could be estimated directly by tracking the daily profits earned on the parity allocations from the NYSE limit order book. Unfortunately, the data needed to produce these estimates, the NYSE's Display Book Data Log Extractor files, are not publicly available.[17] We use an alternate strategy to derive a value for parity trading. Our estimates examine the outcome(s) of each order on the NYSE's public order book that would have fully or partially executed against an incoming marketable order without parity. We can illustrate our approach by looking back to Panel B of Table I. All shares in order 1 fill with and without parity, so it incurs no cost and is not included in our calculations. Order 2 would have filled completely without parity but received only 100 shares with parity; its 100 remaining shares are included in our analysis. Order 3 would have filled 100 shares without parity. We also consider those 100 shares in our analysis. We determine whether orders similar to orders 2 and 3 are ultimately deleted or if they execute at prices that are different than the parity trade price. We refer to orders that receive fewer filled shares with parity than they would have with strict time priority as bypassed orders.

To compute our first estimate of parity's value, we multiply the number of potentially disadvantaged shares in a bypassed order that ultimately executes by the difference between the parity trade price and the eventual trade price. This estimate is a lower bound estimate of the value of parity as it ignores costs associated with unfilled shares in bypassed orders. For this reason, we construct two estimates using the deleted portion of the bypassed orders. First, we compute the opportunity cost associated with bypassed orders by comparing the parity trade price to the end-of-day bid-ask quote mid-point. This difference is signed so that negative numbers indicate an opportunity cost to the bypassed trader (i.e., the price moved in favor of the

[17] See O'Hara, Saar and Zhong (2015) for a description of these data.

parity trader stepping in front of the bypassed trader). Second, we assume that the deleted limit order is replaced with a marketable order that executes at the opposite-side of the quote prevailing when the parity trade occurs. This estimate assumes that the investors whose orders are bypassed are committed to their trade and willing to cross the market by placing a marketable order to complete their trade.

We create an order history for all orders bypassed by orders trading out of time-priority sequence. We report some descriptive statistics on these bypassed orders in Table IV. There are over 39 million liquidity demanding trades and 5.3 million bypassed orders in our sample. Overall, 56% of the 5,321,352 bypassed orders in May 2017 receive at least a partial fill and (not reported in Table IV) 54% receive a complete fill. Finally, 2,420,910 of the bypassed orders are deleted.

For bypassed orders that experience at least one trade after being bypassed, we record the average trade price and compare that price to the price of the trade bypassing the order. Differences between the actual trade price and the price at which the order would have traded if strict time priority was used are signed such that positive numbers are favorable from the bypassed order's perspective and negative numbers unfavorable. That is, positive numbers indicate that the trader buys (sells) at a lower (higher) price than the price associated with the bypassing trade and negative numbers indicate that the trader buys (sells) at a higher (lower) price than the reference price.

The vast majority (95.4%) of bypassed orders eventually trading fill at the same price as they would have filled if the NYSE followed strict time priority. From Table IV, we see that the mean (median) time between being bypassed and the final time stamp of orders receiving a trade is 46.87 (.04) seconds. There are 76,901 orders that fill at worse prices than the reference price

and 57,648 orders that fill at better prices. Multiplying the mean signed difference in price per share times the number of shares traded, we find that public limit order book trades fill at prices that are $56,448.81 worse than they would have filled with strict time priority in May 2017.

Addressing the cost associated with deleted bypassed orders requires an assumption regarding the behavior of the trader who deleted the order.[18] First, we compute realized spreads comparing the price missed to the quote midpoint at the time of the last trade on the day the order was bypassed. This opportunity cost imposed on deleted bypassed orders is $709,063.97. If we assume investors who delete bypassed orders convert to an aggressive order and trade at the price of the opposite-sided quote, we find a cost of $2,904,364.28. Although we recognize that not all deleted limit orders are converted to a marketable order, the latter figure represents an upper bound.

On an annualized basis, our results suggest potential profits to floor-based traders of between $9.2 million (executed bypassed orders and the opportunity cost of deleted bypassed orders) and $35.5 million (executed bypassed orders and assuming that deleted bypassed orders become aggressive) are generated by parity trading in 2017. These profits ignore the revenue DMMs and floor brokers generate using their privileged positions to trade at the open and the close. For example, Bacidore et al. (2013) note that floor brokers and DMMs have several advantages over other market participants in NYSE's closing auction in that they get a "sneak-peak" at the close book every 15 seconds, which includes the MOC/LOC interests, imbalances, and the closing offset interest, which is not available to other market participants.

It is worth noting that our estimates of parity's value are surely lower bounds of the cost imposed on orders affected by parity trades. Specifically, we do not consider orders in the queue that would not have filled for the particular trade in strict time priority but are none-the-less

[18] With non-public NYSE data, one could more accurately determine the actions of traders who delete orders.

potentially disadvantaged. Again, refer to the example in Table I. Order 1 is not disadvantaged by parity. Order 2 has 100 shares that are not disadvantaged by the parity trade of interest because they trade at the same time and price with parity as they would have with time priority and 100 shares that would have traded in the example trade with time priority but do not because of parity. Order 3 would have filled 100 of its 200 shares without parity but received no shares with parity. We consider the 100 shares that would have filled without parity but ignore the 100 shares that would not have filled even with strict time priority. Likewise, we ignore order 4. We do not include these shares in our analysis for two reasons. Computationally, the cascades caused to multiple orders being pushed back in the book multiple times is extremely complex to code. In essence, we would need to create a hypothetical NYSE in which time priority was strictly enforced for all orders. In addition to the complexity of doing this, we would need to make extreme assumptions about trader behavior including assuming that traders would continue to place the same orders in the NYSE book if parity were not part of the NYSE rules.

VIII. Concluding Remarks

To our knowledge, this represents the first paper to address empirically the NYSE's parity rule, which allows deviations from time priority across different types of members. By permitting Designated Market Makers (DMMs) and floor brokers to trade "on parity" with the top of the public limit order book without regard to the time the order is displayed, the NYSE provides a potentially valuable advantage to these participants. The NYSE maintains that their floor-based model provides services to the investing public that improves market quality. Payment for these services is provided by special trading privileges. The value of these privileges is not publicly transparent. We make an effort to quantify the value of one of these

rules. Only after an examination of the value of parity trading can we begin to assess whether the possible benefits of floor traders are worth the apparent costs.

We find that parity trades are observed when it is more difficult to obtain priority in other ways and when it is more likely that an order is in a long, slow-moving queue. Specifically, parity trades are more prevalent when the NYSE quoted spread is at minimum variation, when the stock's price is low, and when alternative methods of obtaining priority on competing exchanges are popular. We also find that parity trades are more likely to occur when the NYSE limit order queue is long and when market and stock price volatility is low.

To estimate parity's cost, we track orders bypassed by parity executions. For bypassed shares that eventually fill, we find that many traders chase price and trade only by paying more (receiving less) than their initial price when they eventually trade. The cost to this subset of traders bypassed exceeds $56,000 in our sample month. We believe that this represents a very conservative lower bound as it ignores the cost associated with deleted bypassed orders. The total costs associated with parity depend on what one assumes that traders do who delete orders when denied fills by parity trades. We consider two cost estimates for deleted bypassed shares. Assuming that the trader bypassed converts to a marketable order after deleting the limit order and reaches across the bid-ask spread leads to a cost of almost $2.9 million in the sample month. A less aggressive estimate of the cost of deleted bypassed orders is to compute the end-of-day opportunity cost of missing out on the trade. This produces a cost of over $709,000. Assuming that May 2017 is typical provides an annual cost parity trading of between $9.2 million and $35.5 million. We recognize that floor brokers have customers so that the value of parity trading (less the floor broker fee) is earned by investors using floor brokers. However, this value still represents a transfer payment from one investor to another. In addition, we are not suggesting

that disadvantaged traders would receive this value as a windfall should the NYSE eliminate parity as the dynamics of order routing decisions are complex. We simply offer these as costs of the current system to shares bypassed by a particular parity allocation.

The potential advantage of floor-based traders is that they provide an alternative source of liquidity to investors and can use their location on the floor to observe market conditions and use judgement in executing orders. DMMs have obligations requiring them to provide liquidity. Floor brokers have no such obligations although extant literature suggests that floor brokers also provide liquidity. Although we cannot separate floor broker and DMM parity trading with our data, it is useful to contrast the magnitude of our numbers to the situation where DMMs are paid directly for their services. In contrast to practice in Europe, listed companies in the U.S. cannot directly contract with DMMs to obtain market making services. As a result, the NYSE uses enhanced rebates and parity to compensate DMMs for subjecting themselves to "rule-based mandatory obligations that carry regulatory risk".[19] Although the NYSE's arms-length compensation of DMMs may reduce the likelihood of market manipulation, it makes it difficult to determine the total liquidity-provision payments. Moreover, rather than coming directly from the listed company, parity revenue is derived by transfer payments between exchange members.

Is the NYSE floor overcompensated relative to what DMMs would receive if they were paid directly by the listed companies for their services? Assume for the moment that parity trading would not exist if it were not a mechanism used by the NYSE to compensate DMMs for providing liquidity. In this situation, we can use our estimates of the costs that parity trading impose on the NYSE's non-floor traders as an estimate of the cost of using trading privileges (rather than direct payments) to compensate the DMM. We focus on parity trading in the 373

[19] More specifically, the DMM must guarantee that all marketable trading interest is satisfied in the opening and closing auctions, must contribute to the NBBO a certain percentage throughout the day, and must provide depth to dampen price moves throughout the day. See https://www.sec.gov/comments/265-29/26529-79.pdf.

NYSE-listed stocks that were part of the S&P500 in May 2017 assuming that parity is more valuable in more actively-traded stocks with institutional trading interest. The results in Table IV indicate parity trading led to between $278,605 and $1,284,391 in profits in May 2017 (about 40% of the total). This suggests parity cost S&P500 investors using the NYSE's public limit order book between $747 and $3,443 per stock during May 2017. In Europe, where firms contract directly with DMMs for the provision of liquidity services, Anand, et al. (2009) find that the most a firm paid a liquidity supplier on the Stockholm Stock Exchange in 2003 was $3,300 per month while Skjeltorp and Odegaard (2015) find issuing firms pay market makers between $4,167 and $8,333 per month on the Nasdaq OMX. In each of these studies, the authors find that the firms that choose to purchase liquidity services tend to be small. Paying for liquidity in highly liquid stocks seems somewhat counter-intuitive. It seems possible that the NYSE's approach of paying for floor liquidity provision via trading privileges is less useful in stocks with low trading volume and requires cross-subsidies from the highly liquid stocks.

Similar to all studies of market structure, our analyses are affected by the endogenous nature of observed order routing practices. It is possible that actual costs exceed our estimate. Sophisticated order routers most likely consider the slower execution of non-floor limit orders routed to the NYSE even if the routers do not explicitly model parity. This might result in different choices of order type and/or order destination than would occur without parity. The costs measured in our paper are those (willingly) incurred by traders choosing to send non-floor limit orders to the NYSE despite the effects of parity. We concede that there might be additional costs associated with traders being more aggressive on the NYSE, routing to venues with slower execution than the NYSE would provide without parity, and/or NYSE limit order traders providing less competitive bid-ask spreads than they would provide without parity. Finally, as

noted earlier, we consider an order as bypassed only if it is denied an execution because of an out-of-time-sequence trade allocation. Orders on the book that would not have received a fill even with strict time priority for a given trade are ignored in our computations. However, these orders might in fact be harmed over a sequence of trades as parity trading pushes them further down the queue.

References

Anand, A., M. Samadi, J. Sokobin, and K. Vendataraman, 2019, Institutional order handling and broker-affiliated trading venues, FINRA Office of the Chief Economist, Working Paper.

Anand, A., and K. Venkataraman, 2016, Market conditions, fragility and the economics of market making, *Journal of Financial Economics* 121, 327-349

Anand, A., C. Tanggaard, and D. Weaver, 2009, Paying for market quality, *Journal of Financial and Quantitative Analysis* 44, 1427-1457.

Baccidore, J., B. Polidore, and W. Xu, 2013, The cost of the d-quote, *ITG Blotter* 4, October 23.

Battalio, R., A. Ellul, and R. Jennings, 2007, Reputation effects in trading on the NYSE, *Journal of Finance* 62, 1243-1271.

Battalio, R., S. Corwin, and R. Jennings, 2016, Can brokers have it all? On the relation between make-take fees and limit order execution quality, *Journal of Finance* 71.

Bessembinder, H., A. Carrion, L. Tuttle, and K. Venkataraman, 2016, Liquidity, resiliency and market quality around predictable events: Theory and evidence, *Journal of Financial Economics*.

Bessembinder, H., J. Hao, and M. Lemmon, 2011, Why designate market makers? Affirmative obligations and market quality, Working Paper, Arizona State University.

Bessembinder, H., J. Hao, and Kuncheng Zheng, 2017, Liquidity Provision Contracts and Market Quality: Evidence from the New York Stock Exchange, Working Paper, Arizona State University.

Clark-Joseph, A., M. Ye, and C. Zi, 2017, Designated market makers still matter: Evidence from two natural experiments, *Journal of Financial Economics* 126, 652-667.

Cordella, T., and T. Foucault 1999, Minimum price variations, time priority and quote dynamics, *Journal of Financial Intermediation* 8, 141-173.

Frino, A., A. Hill, and E. Jarnecic, 2000, An empirical analysis of price and time priority and pro rata trade execution algorithms in screen-traded markets, *Journal of Derivatives* 7, 41-48.

Gwilyn, O., and E. Alibo, 2003, Decreased price clustering in FTSE100 futures contracts following a transfer from floor to electronic trading, *Journal of Futures Markets* 23, 647-659.

Glosten, L., 1994, Is the electronic open limit order book inevitable?, *Journal of Finance* 49, 1127-1161.

Harris, L., 1998, Optimal dynamic order submission strategies in some stylized trading problems, *Financial Institutions and Instruments* 7, 1-76.

Harris, L., 2003, Trading and Exchanges: Market Microstructure for Practitioners, Oxford University Press.

Harris, L., and J. Hasbrouck, 1996, Market versus limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213-232.

Harris, L., and V. Panchapagesan, 2005, The information content of the limit order book: Evidence from NYSE specialist trading decisions, *Journal of Financial Markets* 8, 25-67.

Hasbrouck, J., 2007, Empirical Market Microstructure, Oxford University Press.

Ikenberry, D., and J. Weston, 2008, Clustering in US stock prices after decimalization," *European Financial Management* 14, 30-54.

Lee, C. and M. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733-746.

Moallemi, C., and K. Yuan, 2017, A model for queue position valuation in a limit order book, Columbia Business School Working Paper 17-70.

O'Hara, M., G. Saar, and Z. Zhong, 2015, Relative tick size and the trading environment, *Review of Asset Pricing*.

Panchapagesan, V., 1997, What if time priority is not enforced among traders?, University of Southern California working paper.

Panayides, M., 2007, Affirmative obligations and market making with inventory, *Journal of Financial Economics* 86, 513-542.

Pirrong, C., 1996, Market liquidity and depth in computerized and open outcry trading systems," *Journal of Futures Markets* 16, 519-543.

Ready, M., 1999, The specialist's discretion: Stopped orders and price improvement, *Review of Financial Studies* 12, 1075-1112.

Skjeltorp, J., and B Odegaard, 2014, When do firms pay for market making in their own stock?, *Financial Management* 44.

Werner, I., and G. Sofianos, 2000, The trades of NYSE floor brokers, *Journal of Financial Markets* 3, 139-176.

Table I. Hypothetical Parity Example

All orders in the table share the same limit price. They differ by when they were displayed on the book and order numbers reflect the ordinal time of order arrival. The "trades" enumerated in Panel B are for illustrative purposes only and do not represent the prints that would occur in actual trade reporting.

Panel A – Orders seeking to buy shares at the NBB

Order Number Based on Time Priority	Order Size (Shares)	Order Category
1	100	Public Limit Order
2	200	Public Limit Order
3	200	Public Limit Order
4	100	Public Limit Order
5	100	Floor Broker
6	100	Designated Market Maker

Panel B – Trades Associated with a 400 Share Marketable Sell Order.

Order 1		Order 2		Order 3		Order 4		Order 5		Traded	
Trade	Leave	Trade	Leave	Trade	Leave	Trade	Leave	Trade	Leave	Trade	Cum
100	0					100	0	100	0	300	300
		100	100							100	400

Table II. Descriptive statistics of all trades and trades involving at least one parity allocation executed on the NYSE in May 2017.

Statistics in this table require the trade to occur between 9:45 and 15:45 and that the trade price is between $1 and $1000. With three exceptions, the variables in the table below are computed taking the equally-weighted average for the trades during the month. Price Volatility is computed by taking the standard deviation of trade prices in a given stock. Number of Trades and Number of Shares Traded are the total number of trades and shares traded for a given stock during the month. Statistics for these three variables are the stock-weighted average, minimum, and maximum.

Panel A. All NYSE trades. N = 38,895,277 trades.

	Mean	Minimum	Maximum
Trade Price	$51.89	$1.01	$990.16
Number of Trades	43,116	1	200,630
Number of Shares Traded	11,488,297	50	186,523,397
Price Volatility	1.29	0	43.63
Depth (shares)	978	1	1,005,179
% of Trades Executed when NYSE Bid Ask Spread is a Penny	66.83%	0	100%
Trade Time	12:35:32	9:45:00	15:45:00
# of Orders Bypassed	0.18	0	202
# of Orders Participating in Trade	1.77	1	756
% of Trades with a Parity Allocation	14.69%	0	100%
% of Shares Allocated to Parity Orders	9.89%	0	100%
Trade Size	251	1	3,812,600

Panel B. NYSE trades with and without a parity allocation. There are 5,724,506 trades with at least one parity allocation.

	Mean	
	Trades with a Parity Allocation	Trades without a Parity Allocation
Trade Price	$33.83	$55.02
Number of Trades	46,924	42,459
Number of Shares Traded	20,610,151	10,032,106
Depth (shares)	3,512	517
% of Trades Executed when NYSE Bid Ask Spread is a Penny	87.42%	63.28%
# of Orders Participating in Trade	2.43	1.66
% of Shares Allocated to Parity Orders	67.37%	0.00%
Trade Size	443	218

Table III. Multivariate analysis of the % of shares in a trade allocated to parity orders.

We use Ordinary Least Squares where the dependent variable is percent of a marketable order allocated to orders trading at parity on the NYSE. We use a logistic regression where the dependent variable is a binary variable equal to 1 if the order contains at least one parity order. We examine trades executed on the NYSE in May 2017 between 9:45:00 and 15:45:00 with execution prices between $1.00 and $1,000.00. *Penny Spread* equals one if the execution-time NYSE quoted spread equals $1.00 and 0 otherwise. *Relevant-Side Depth* is the number of shares offered (bid for) at the NYSE best offer (bid) for trades initiated by the buyer. For trades initiated by the buyer (seller), *New Passively Set NYSE Price* equals 1 if the execution-time NYSE offer (bid) price is higher (lower) than the NYSE offer (bid) price that was prevailing when the prior trade in the stock was executed. *Seconds Past Midnight* is the trade's execution time in seconds past midnight. *Absolute CRSP Market Return* is the absolute return on the day in which the trade occurs. *Seconds Past Midnight* is the execution time of the trade in seconds past midnight. *Ln (Shares Traded on the NYSE)* is the natural log of the number of shares in sample trades in the stock executed on the NYSE in May 2017. *Average Price* is the average price of trades in the stock executed on the NYSE in May 2017. *Trade Price Volatility* is the standard deviation of the prices of the trades in the stock executed by the NYSE in May 2017. *Low Fee Exchanges' Market Share* is the percentage of trades in a given stock in the last seven trading days of May 2017 that are executed by the three exchanges with lower take fees than the NYSE.

	Parameter Estimate OLS	Parameter Estimate Logistic
Intercept	-0.0619	-7.0924
Penny Spread	0.0534	0.5162
Relevant-Side Depth	6.28×10^{-6}	1.2×10^{-4}
New Passively Set NYSE Price	0.0408	0.4151
Absolute CRSP Market Return	-0.4596	-3.6966
Seconds Past Midnight	1.47×10^{-6}	2.0×10^{-5}
Active Side of Trade is a Seller	0.0061	0.0635
Trade is Between 15:30:00 & 15:45:00	0.0162	0.2114
Ln (Shares Traded on the NYSE)	1.27×10^{-9}	0.2571
Average Price	-2.79×10^{-4}	-0.00664
Trade Price Volatility	-0.0027	-0.2275
Low Fee Exchanges' Market Share	0.3828	1.9529
N = 38,292,542 trades R^2 = 7.72%		

Table IV

We report statics summarizing trading on the NYSE during May 2017. Liquidity-demanding trades is all trades during the sample month.
An order is defined to be bypassed if the order would have filled completely if one or more parity orders had not stepped ahead of the order.

	Full Sample	SP500 Stocks
Liquidity demanding trades	39,029,914	16,796,164
Bypassed orders	5,321,352	2,162,446
Bypassed orders with an eventual fill	2,961,420	1,171,164
Number of shares executed in bypassed orders	308,984,748	108,326,814
Mean (median) time elapsing between the time the order fills and the time of the parity trade – in seconds	46.87 (0.04)	20.55 (0.007)
Bypassed orders that eventually fill at a worse price than the parity price	76,901	36,938
Bypassed orders that eventually fill at the same price as the parity price	2,826,871	1,108,234
Bypassed orders that eventually fill at a better price than the parity price	57,648	26,119
Mean share-weighted difference in order and parity trade price for bypassed orders	-$0.00024	-0.00030
Bypassed orders that are deleted without a complete fill	2,420,910	1,014,213
Number of bypassed shares in bypassed orders that are deleted without a complete fill	328,176,918	115,985,399

Data Appendix.

A. Individual message counts.

Appendix Table I provides counts of the messages of interest during the sample month of May 2017. For reference, these messages originate from an average of 3,181 traded securities per sample day. There are nearly 3.4 billion messages: 43.6% of the messages are new orders, 41.6% are deletions, 9% are order replacements, 3.1% are order modifications and 2.7% are executions. Traders delete over 95% of orders placed.

B.1 Descriptive statistics on messages.

Appendix Table II reports select descriptive statistics on messages. We report the statistics on four samples – the full sample, the sample of messages with a time stamp between 9:45 and 15:45 and analogous samples with a price screen to eliminate the effect of very low-priced and very high-priced stocks. Starting at 9:45am abstracts away from the highly unusual opening trade. Ending at 15:45 excludes the Closing Imbalance Period, which restricts traders' flexibility with some order types. Stocks priced below $1 trade on a finer tick size than those priced at $1 or more and stocks like Berkshire Hathaway A shares, with a share price exceeding $250,000 during the sample period, distort statistical analyses.

The add messages' statistics are presented in the first row of each panel in Appendix Table II. Initially, when examining all 1,482,074,696 new orders, we find that they arrive about mid-day on average. The average order size is 299 shares and there are almost an equal number of buy and sell orders. Requiring that the orders arrive between 9:45 and 15:45 reduces the sample size by about 15%, but makes little difference in the sample mean arrival time or order size or mix of buy and sell orders. Focusing our attention on add messages with limit prices

between $1 and $1000 eliminates about 0.4% of the sample. The price restriction on the sample has almost no effect on the average order arrival time or size or order side.

We begin with 92.9 million execution messages. The average time stamp on executions is the latest among the four message types examined. When we restrict the analysis to executions that happen in the reduced trading day, the number of observations drops by more than 25% from the total to 69,181,963 (many orders execute at open and in the final 15 minutes of the trading day). Ignoring the executions with prices less than $1 and more than $1000 eliminates only 0.3% of the sample. Fills occur, on average, roughly at mid-day. The average price of a filled order is $53-$54 ($44-$45 if we exclude extremely priced shares) for about 175 shares (fills within the 9:45-15:45 period are somewhat smaller, 142 shares).

We reduce the 105 million modification messages by about 11.7% when focusing our attention on the 9:45-15:45 period. Eliminating the extreme stock price trades discards about 0.2% of the data. Modifications, although still occurring at roughly mid-day, occur slightly earlier than executions, on average. The average modified order size exceeds 500 shares, significantly larger than the number of shares in add or fill messages.

We calculate descriptive statistics for the 304,457,510 replace messages and report them in the final row of each Panel in Appendix Table II. The descriptive statistics for the replace messages are similar to the add messages, although order replacements seem to be directed toward orders that are slightly larger than the average new order. Very few replace messages occur in extreme priced stocks – the price screen eliminates only about 0.05% of the replace sample. Again, not surprisingly, order replacements occur prior to order executions on average.

B.2 Descriptive statistics on order history.

Each of the roughly 1.5 billion Add Messages during May 2017 in Appendix Table I represents a new order. Each order builds a history; a trader can modify/replace (possibly multiple times) the initial order, the order can receive one or more executions, and/or a trader can delete the order before it fills. We create each new order's entire history by linking all modify, replace, execute, and delete messages to the original order. For some orders, this involves literally thousands of messages.

Appendix Table III provides descriptive statistics regarding order history (regardless of intraday timing and/or price level). Panel A summarizes the data on all new orders. Average initial order size is just under 300 shares, but there are very small and extremely large orders. Although the average order has few modifications, replacements, and/or executions, some orders have a very complex history. Of the new orders represented by Add Messages, 96.8% end (with a delete, a complete fill, or reach the end of the trading day) with their original order identification number, i.e., the trader chooses not to replace them. Only slightly fewer orders end with their original price (93.94%) or size (92.86%) – results not reported in Appendix Table III.

From Panel F, we see that the vast majority of new orders, about 95%, end with a deletion. Deleted orders do not seem remarkably different in size or price nor in the frequency of modifications/replacements from the average new order. Consistent with the fact that the order is cancelled, there are fewer trades on deleted orders that the average new order. The mean (median) time elapsing between order arrival and deletion is about 7.5 minutes (three seconds). The skewness in the order life comes from orders that arrive early in the trading day and sit on the order book for much of the trading day before deletion.

Orders can be modified (retaining the original order identification number) or cancelled and replaced (creating a new order identification number). Panels B and C of Appendix Table III report that orders with modify (replace) messages represent about 5.5% (3.2%) of new orders yet, on average, there are .07 modifications (.21 replacements) per new order. This is consistent with the fact that traders modify (replace) some orders frequently.[20] The maximum number of modifications (replacements) is 13,402 (8,910), which represents a modification (replacement) every 0.003 (1.01) seconds of that order's life. Thus, traders monitor some orders quite closely and most likely via an automated order algorithm. From Panel D we find that about 8.4% of new orders receive either a modification or a replacement message, however only 0.27% of orders receive both types of messages. Orders with modifications/replacements initially are larger than the overall average order. The ending size after being modified/replaced/executed is similar to the overall average order size. Orders that are modified/replaced average roughly twice the number of trades compared to the typical initial order suggesting that the changes in order specifications improve the chances of a trade. In Panel E, we see that trades occur in 5.2% of orders. Some orders experience many executions; the maximum number of executions is 1,176. The fact that the mean execution price ($53.58) is less than the mean new order price ($158.10) suggests that there are many more executions of lower priced securities than of higher priced stocks.

Panel B of Appendix Table III contains statistics on orders with histories that contain a modification message. Conditional on having at least one modification, the average number of modifications is 1.29. Orders that traders modify have larger initial sizes, on average, than the typical new order. The ending order size, affected by both modifications and executions, is comparable to the initial new order size. Order price changes only slightly, on average, over the

[20] It would be nice if we can talk about e-quotes used by floor brokers are the orders that are cancelled and replaced.

life of modified orders. In Panel C, we summarize data on orders with a cancel-replace message(s) in their lives. Conditional on having at least one cancel-replace message, the average exceeds six replacements. Orders with replacements seem to be different from orders with modifications. The initial order size of orders with replacement messages is smaller than for orders with modifications and the average price is much smaller. The cancel-replace intensity seems to be higher than the approach of modifying orders (conditional mean of 6.3 replacements compared to 1.3 modifications). Orders that receive a trade on average receive 1.21 trades (Panel E) and leave 22 shares unfilled.

Traders can change either the order's price, the order's size, or both during the order's life. Appendix Table IV provides descriptive statistics on these order changes. About 6.45% of orders receive a price change and 2.18% receive a volume change. On average, orders that traders modify and/or replace receive .4785 (.0715 increases and .4070 decreases) size changes and 2.8057 (1.4046 increases and 1.4011 decreases) price changes. It is nearly equally likely that order's price is increased or decreased. It is much more likely that traders decrease the order's original size than they increase size (possibly, because they lose time priority with a size increase). Conditional on a modification or replacement that alters price (size), the mean number of price (size) changes is 3.67 (1.86) and the maximum number of price (size) changes is 6,704 (3,641). Traders seem to chase price. For buy orders, conditional on a price change from the original order, there are an average of 2.25 price increases and 1.51 price decreases per changed order. For sell orders with a price change, there are an average of 1.34 price increases and 2.06 price decreases. Thus, it is approximately 50% more likely that traders changing price adjust the order to a less favorable price than the initial price target.

Also reported in Appendix Table IV are descriptive statistics on the time between all messages in orders that traders modify and/or replace. This includes modify, replace, trade, and delete messages. Although the mean time between messages is approximately one minute, orders that stay on the book from pre-open to market close exert an undue influence on the mean time. The median time is about 0.7 seconds. Thus, traders manage the typical modified/replaced order quite actively.

Finally, Appendix Table V provides additional detail on orders that receive executions. Of the 76,472,791 orders that receive a trade, 92% fill completely and about 76% fill completely without any changes to their original terms. Orders that fill completely have slightly smaller order sizes than the average new order or the average order receiving a trade.[21]

[21] We can focus on activity away from the market open and close and/or in securities with prices in the $1 to $1000 range. Eliminating orders outside the 9:45-15:45 window reduces the sample size by about 15%. However, differences from the overall results are minor. There is a slight increase in the frequency of modifications and a somewhat larger decrease in the frequency of replacement orders; the net effect is that changed orders in the restricted sample are about 7.2% of the total compared to 8.4% in the overall sample. It also is the case that there are fewer trades in the clock-constrained sample; 4.2% of the orders experience a trade compared to 5.4% overall. This is likely because the final 15 minutes of trading contains a disproportionally large amount of trading. Not surprisingly, the mean (but not median) time between events falls substantially when we eliminate orders arriving prior to market open. Other than a very large reduction in mean prices, eliminating the 0.5% of the sample with extremely low and extremely high prices makes very little difference in the reported results.

Appendix Table I. Message counts by message type.

Date	Records	Message Type					Market Return*
		Add	Modify	Delete	Execute	Replace	
20170501	143,907,450	62,532,536	4,024,320	59,421,317	4,082,836	13,846,441	0.001959
20170502	154,768,102	67,597,382	4,941,263	64,077,247	4,566,993	13,585,217	0.000241
20170503	162,386,848	70,731,409	5,017,664	67,183,798	4,672,840	14,781,137	-0.002763
20170504	175,093,606	76,283,594	5,402,388	72,349,869	5,128,474	15,929,281	-0.000856
20170505	152,462,647	66,448,320	4,876,333	63,226,176	4,230,875	13,680,943	0.005768
20170508	138,302,854	60,526,381	4,132,064	57,503,846	4,009,881	12,130,682	-0.001087
20170509	149,490,122	65,100,823	4,427,052	61,829,634	4,291,804	13,840,809	-0.001032
20170510	151,478,803	65,916,943	4,623,822	62,689,896	4,211,598	14,036,544	0.002312
20170511	157,523,367	68,455,621	5,009,828	65,191,860	4,240,495	14,625,563	-0.003002
20170512	137,472,459	59,835,898	4,302,837	56,957,229	3,794,063	12,582,432	-0.001566
20170515	134,789,592	58,781,522	4,212,975	55,850,599	3,877,968	12,066,528	0.00551
20170516	141,123,053	61,543,256	4,361,728	58,577,468	3,915,093	12,725,508	-0.000537
20170517	206,188,968	89,936,920	6,086,715	86,079,708	5,012,183	19,073,442	-0.018313
20170518	215,705,872	93,944,726	6,619,690	90,085,683	5,040,077	20,015,696	0.003335
20170519	164,575,810	71,805,574	5,210,939	68,414,887	4,433,613	14,710,797	0.007539
20170522	138,923,867	60,916,843	4,066,989	58,008,080	3,855,014	12,076,941	0.005368
20170523	144,161,894	63,024,880	4,533,116	60,122,019	3,847,242	12,634,637	0.001515
20170524	147,036,498	64,350,122	4,799,766	61,407,636	3,860,346	12,618,628	0.002374
20170525	155,410,057	68,004,741	5,086,535	64,820,288	4,167,993	13,330,500	0.003144
20170526	124,570,155	54,515,661	3,846,406	51,960,277	3,402,151	10,845,660	0.000189
20170530	129,688,855	56,849,868	3,960,099	54,110,963	3,645,301	11,122,624	-0.002397
20170531	170,809,918	74,971,676	5,514,513	71,473,031	4,653,198	14,197,500	-0.000444
Totals	**3,395,870,797**	**1,482,074,696**	**105,057,042**	**1,411,341,511**	**92,940,038**	**304,457,510**	

*CRSP Value-weighted market return without distributions.

Appendix Table II. Descriptive Statistics – All Message Types NYSE TAQ Integrated Feed, May 2017

Panel A – All Observations.

Message Type	# Observations	Average			% Buy
		Arrival Time	Trade Price	Volume (shrs.)	
Add	1,482,074,696	12:28:19	n.a.	299	50.05
Execute	92,940,038	13:01:10	$53.87	175	n.a.
Modify	105,057,042	12:29:00	n.a.	551	n.a.
Replace	304,457,510	12:28:29	n.a.	318	n.a.

Panel B – Observations within 9:45:00 and 15:45:00.

Message Type	# Observations	Average			% Buy
		Arrival Time	Trade Price	Volume (shrs.)	
Add	1,257,901,204	12:28:19	n.a.	268	49.99
Execute	69,181,963	12:38:36	$53.72	142	n.a.
Modify	92,724,376	12:23:49	n.a.	523	n.a.
Replace	262,113,670	12:21:32	n.a.	296	n.a.

Panel C – Observations with $1.00 $\leq$ price $\leq$ $1,000.00.

Message Type	# Observations	Average			% Buy
		Arrival Time	Trade Price	Volume (shrs.)	
Add	1,476,047,494	12:28:07	n.a.	294	50.04
Execute	92,693,161	13:01:14	$45.30	174	n.a.
Modify	104,894,325	12:29:00	n.a.	549	n.a.
Replace	304,318,219	12:28:27	n.a.	317	n.a.

Panel D – Observations with $1.00 $\leq$ price $\leq$ $1,000.00 and within 9:45:00 and 15:45:00.

Message Type	# Observations	Average			% Buy
		Arrival Time	Trade Price	Volume (shrs.)	
Add	1,255,260,525	12:27:28	n.a.	266	49.99
Execute	68,988,158	12:38:37	$43.68	142	n.a.
Modify	92,586,779	12:23:49	n.a.	521	n.a.
Replace	262,039,498	12:21:32	n.a.	296	n.a.

n.a. = Not Applicable

Appendix Table III. Basic descriptive statistics for sample orders.

Panel A. Add messages (N = 1,481,792,735)

	Mean	Minimum	Maximum
Original Volume (shares)	299	1	4,000,000
Original Price	$158.10	$0.0001	$347,300.00
Number of Modifications	0.0709	0	13,402
Number of Replacements	0.2052	0	8,910
Number of Trades	0.0630	0	1,176

Panel B. Orders with modify messages (N = 81,233,917)

	Mean	Minimum	Maximum
Original Volume (shares)	536	1	500,000
Ending Volume (shares)	279	0	2,500,000
Original Price	$101.23	$0.0012	$261,970.00
Ending Price	$100.99	$0.0012	$261,970.00
Number of Modifications	1.2919	1	13,402
Number of Replacements	0.3720	0	2,877
Number of Trades	0.1320	0	190

Panel C. Orders with replace messages (N = 47,887,411)

	Mean	Minimum	Maximum
Original Volume (shares)	339	1	990,000
Ending Volume (shares)	313	0	990,000
Original Price	$50.17	$0.0010	$262,000.00
Ending Price	$50.17	$0.0010	$260,000.00
Number of Modifications	0.1631	0	6,076
Number of Replacements	6.3450	1	8,910
Number of Trades	0.1255	0	376

Panel D. Orders with modify or replace messages (N = 125,092,314)

	Mean	Minimum	Maximum
Original Volume (shares)	453	1	990,000
Ending Volume (shares)	280	0	2,500,000
Original Price	$84.02	$0.0010	$262,000.00
Ending Price	$83.86	$0.0010	$261,970.00
Number of Modifications	0.8389	0	13,402
Number of Replacements	2.4295	0	8,910
Number of Trades	0.1273	0	376

Appendix Table III (continued)

Panel E. Orders with trade messages (N = 76,472,791)

	Mean	Minimum	Maximum
Original Volume (shares)	242	1	3,823,300
Ending Volume (shares)	22	0	988,734
Execution Price	$53.58	$0.0010	$252,000.00
Number of Modifications	0.1560	0	13,402
Number of Replacements	0.4487	0	7,533
Number of Trades	1.2152	1	1,176

Panel F. Orders with delete messages (N = 1,441,502)

	Mean	Minimum	Maximum
Original Volume (shares)	303	1	4,000,000
Ending Volume (shares)	289	1	4,000,000
Original Price	$158.42	$0.0001	$347,300.00
Ending Price	$158.41	$0.0001	$347,300.00
Number of Modifications	0.0667	0	4,154
Number of Replacements	0.1937	0	8,910
Number of Trades	0.0053	0	274
Seconds from Add to Delete	450.61	0.000005	36,899.0

Appendix Table IV. Descriptive statistics for orders that are modified and/or replaced.

Panel A. All orders with modify or replace messages (N = 125,092,323)

	Mean	Minimum	Maximum
Number of volume increases	0.0715	0	3,641
Number of volume decreases	0.4070	0	3,638
Number of price increases	1.4046	0	6,698
Number of price decreases	1.4011	0	6,704
Seconds between messages	57.894	0.000004	30,599

Panel B. Buy orders with modify or replace messages (N = 63,083,105)

	Mean	Minimum	Maximum
Number of volume increases	0.0700	0	3,641
Number of volume decreases	0.3987	0	3,638
Number of price increases	1.6972	0	6,002
Number of price decreases	1.1822	0	6,004
Seconds between messages	59.082	0.000004	30,595

Panel C. Sell orders with modify or replace messages (N = 62,009,218)

	Mean	Minimum	Maximum
Number of volume increases	0.0733	0	2,077
Number of volume decreases	0.4156	0	2,077
Number of price increases	1.0718	0	6,698
Number of price decreases	1.6383	0	6,704
Seconds between messages	56.596	0.000005	30,599

Appendix Table V. Descriptive statistics for orders with trade message(s).

Panel A. All orders with modify or replace messages (N = 125,092,323)

	Mean	Minimum	Maximum
Original volume (shares)	242	1	3,823,300
Ending volume (shares)	22	0	988,734
Execution price	$53.58	$0.0010	$252,000.00
Number of trades	1.2152	1	1,176

Panel B. Orders that fill completely (N = 70,591,634)

	Mean	Minimum	Maximum
Original volume (shares)	224	1	3,823,300
Execution price	$53.83	$0.0010	$252,000.00
Number of trades	1.2119	1	1,176

Panel C. Orders that fill completely without modify and/or replace messages (N = 58,392,300)

	Mean	Minimum	Maximum
Original volume (shares)	218	1	3,823,300
Execution price	$57.15	$0.0010	$252,000.00
Number of trades	1.2215	1	1,176