

Competition and Exchange Data Fees

Jonathan Brogaard and James Brugler*

October 2, 2020

Abstract

Exchanges are monopolist suppliers of their own order book data. We examine three events where exchanges begin charging a fee for order book data for the first time and test whether or not these fees affect their market share in a difference-in-differences setting. We find that the introduction of fees leads to a fall of market share of around 5-8 percent. Examining average trading costs, price impact and dealer revenue per trade around the events indicates that order routing decisions of informed traders are relatively more sensitive to order book data fees than other trader categories.

JEL Codes: G12, G14

Keywords: Data Fees, Exchange Competition

* Brogaard is at the David Eccles School of Business, University of Utah, Email: brogaardj@eccles.utah.edu; Phone: +1 (801) 581-7463. Brugler is at the University of Melbourne, Email: james.brugler@unimelb.edu.au; Phone: +61 3 8344 1164.

1 Introduction

A major source of revenue for exchanges is the sale of their data. For example NASDAQ in the 4th quarter of 2018 generated \$97 million in revenues from selling its data.¹ The data that they sell is to some degree a public good (Alan and Schwartz, 2013). It plays an important role in the price discovery process. The SEC is responsible for approving data fee charges that the exchange request. Recently, a spat has erupted between the SEC (and many investors) and the exchanges over whether data fees are too high. This article tests whether investors change their behavior due to fee changes, and how the changes impact market quality.

The exchange data environment and fee structure is complex, and the literature on competitive markets is voluminous.² However, the opposing theories can be summed up concisely. The Managed Funds Association and The Alternative Investment Management Association filled a petition regarding the SEC ruling on August 22, 2018 that sums up the concern: “The Associations are concerned that exchanges as exclusive processors are charging unreasonable fees for market data products, and as a consequence, restricting trade and harming competition.” (MFA and AIMA, 2018). Investors argue that exchanges are exclusive processors of essential information and therefore high data fees are a way to extract economic rents from a captured audience. Exchanges argue that competition amongst the exchanges means investors choose to not subscribe to their data feeds and focus on trading on other venues.

One simple means of examining the claims is to study how market share reacts to data

¹<http://ir.nasdaq.com/news-releases/news-release-details/nasdaq-reports-fourth-quarter-and-full-year-2018-results>

²For example, see https://www.nasdaqtrader.com/Trader.aspx?id=Global_Market_Data or <https://www.nyse.com/market-data> for a list of the products offered by NASDAQ and NYSE, respectively. In addition to the data itself other services, such as colocation, are also available from many of the exchanges.

fee changes. An increase in fees corresponding to a decrease in market share is evidence that exchanges are subject to market forces to some degree. We study three discrete data fee events in the U.S. Equity markets that occurred on ARCA, Direct Edge, and BATS. We find that when a trading venue increases its data fees, it loses market share. The results are consistent with investors being able to pick and choose their data needs and being able to vote with their wallets.

Our results suggest that trading venues must optimize a trade off between increasing data fees and losing investors subscriptions. The results show that some investors choose not to subscribe (i.e. some investors are sensitive to price at the extensive margin).

We next turn to understanding who, at the margin, is changing their trading behavior. While we cannot directly observe the behavior of different types of investors, we can infer whose behavior changed based on changes in market quality. We more fully explore the theoretical mechanisms that would link different types of traders to different market quality outcomes. We focus on three standard categories of market participants: market makers, informed traders, and liquidity traders.

Theory suggests that with less information from a given venue, an investor faces greater execution risk when sending an order to that venue (because they have less information with which to assess execution probability) and also face greater information asymmetry against the traders who do subscribe. Therefore, it is more likely for an investor who does not subscribe to a venues' data feed to route their orders elsewhere. Different categories of traders' orders will have different characteristics. By examining the average characteristics of trades on venues introducing fees compared with other venues, we can infer which category of traders are more sensitive to market data fees.

Using the same setup as we used to examine change in volume, we find that the intro-

duction of a data fee results in price impact falling and realized spread increasing. These changes are most consistent with informed traders decreasing their trading on the venue increasing a data fee. Note that we can only make statements about the relative change among the market participants. It could be, for example, that the introduction of a data fee impacts the trading activity of both market makers and informed traders, but that the informed traders are impacted to a greater degree.

As an alternative approach, if informed traders are the market participant most impacted by the data fee introduction, then we expect to see the largest change in volumes and realized spreads to occur in high price impact stocks. Accordingly, we separate stocks into high and low price impact stocks. As expected if informed traders are departing the market, the high price impact stocks are the ones most impacted.

Our paper adds to the literature on exchange data fees. A broader literature examines other dimensions of exchange competition. For example, Ramos and von Thadden (2008) focuses on how transaction costs impacts exchange competition and Pagnotta and Philippon (2018) studies the role of speed among exchanges. Hendershott and Jones (2005) study the removal of market data for three ETFs trading on Island ECN. Like us, they document a large and significant drop volume traded for these ETFs on Island ECN after the venue ceases to display order book data to any traders for these securities. An important difference between this event and our study is that, in their case, order book information is unavailable at any price to any trader. In our study, market data is available for a price, and so both access to market data and subsequent order routing decisions are strategic decisions. Hendershott and Jones (2005) find that liquidity traders are less likely to route their orders to Island ECN after market data is removed, while we find that informed traders are more sensitive to the existence of fees.

Our focus on data fees specifically most closely relates to Jones (2018) and Hendershott, Rysman, and Schwabe (2020). Jones (2018) argues that though market data is valuable to market participants, market data fees are a small cost to the industry overall and the market for data is characterized by “robust competition”. Hendershott et al. (2020) shows that the NYSE gains market share after the introduction of new market data products, the NYSE Integrated Feed, demonstrating that this information is valuable to traders. They further demonstrate the existence of a positive externality for market data: traders who do not subscribe also increase their trading on the NYSE after the Integrated Feed begins. Our results are consistent with these findings. We contribute to this literature by first documenting that the introduction of fees for market data leads to lower market share, and by identifying informed traders as the most affected trader categories after fees are introduced.

2 Three Fee Introductions

Fees for order book information are complex and frequently changing. Due to the complexity of fees, different market participants face different total costs for market data and it is difficult to summarize the average cost of market data in a simple way. We study three events where exchanges that introduce fees for the first time. These events are advantageous because we can clearly identify instances where the cost of market data unambiguously increase for all market participants.

The first event is the introduction of market data fees by ARCA on 1 January 2009. Prior to this date, order book information from ARCA was provided free of charge. In 2006, the exchange filed with the SEC to establish fees for its market data (SEC, 2008). The filing was subsequently approved in December 2008 and fees were introduced in 2009. These initial

fees included a \$750 per month access fee, a monthly fee of \$15 per device for professional subscribers for ETFs and equities data as governed by the CTA Plan and \$15 per device and month for data covering other equities securities governed by the Nasdaq UTP plan. For non-professional subscribers, the monthly fees per device were \$5 for both categories.

The second event is the introduction of fees for depth of order book data for trading on the Direct Edge EDGX exchange. This real-time data contains information on all displayed orders, executions and order modifications and was initially provided free-of-charge to all Member firms. On 1 May 2012, Direct Edge introduced monthly fees of \$500 for internal use of the feed and \$2,500 for users who wish to distribute the feed externally (CBOE, 2012).

The final event is the introduction of market data fees by BATS exchange on 1 July 2013. BATS fee structure initially comprised of a total charge of \$1,500 per month for depth of book data for internal use only, covering both the BZX and BYZ exchange (Reuters, 2013). For top of book and last sale data, these fees total \$1,000 across both sub-exchanges. For firms that wish to distribute the data externally, the total monthly costs would be \$7,500 for depth of book data and \$5,000 for top of book and last sale data.

3 Hypothesis Development

Traders faced with fees for order book information must decide whether or not the expected benefit of having this information outweighs the explicit costs of purchasing it. It may be that order book information from venues introducing fees is so important that practically all market participants are forced to subscribe to the data after fees are introduced. If so, we argue there should be no change in market conditions (market shares, trading costs across venues etc) after fees are introduced compared with prior to the introduction of fees. Nothing has changed in terms of the amount or distribution of information throughout the trading

economy. Order submission strategies that were optimal in the absence of fees would remain optimal when fees are introduced. In this case, the venue introducing fees suffers no loss of market share and fees are just a mechanism to transfer surplus from the trading process from traders to the exchange.

If instead, at least some traders are sensitive to the introduction of fees at the extensive margin (i.e. they choose not to subscribe to order book data), then we argue that they should become less likely to submit orders to that venue and the venue introducing fees will lose market share. This leads to the first testable hypothesis regarding data fees and trading outcomes: at least some traders are sensitive to the cost of order book information and choose to route their orders elsewhere after data fees are introduced, leading to lower market share.

We test this by examining volume traded by stock before and after the introduction of fees and compare this to the volume traded by stock on other venues with no change in fees. We also view this analysis as a test of the degree of market power that exchanges have in charging fees for order book information. If order book information is sufficiently important to all trader types that they all subscribe, then exchanges appear to have a high degree of market power and there is a high potential for economically inefficient outcomes due to exchanges charging socially inefficient prices for data. If the evidence instead suggests that at least some traders choose not to subscribe, and market share falls, then concerns regarding market power and efficiency are lessened.

We then consider the incentives of different trader types to subscribe to fees and how they may rationally respond when fees are introduced.

The expected benefit of having order book information can differ for different trader types. For example, market makers providing liquidity without this information face an

important source of information asymmetry compared with other market makers who do have this information. It is highly likely that these market makers will face severe adverse selection problems from not being able to update their quotes after changes in the state of the order book. Though order book information would appear to be a key requirement for any market maker wishing to provide liquidity on a particular venue, it does not follow that all market makers would necessarily subscribe to the order book information for a given venue. Instead, smaller market makers who make low levels of profit on the venue may rationally not subscribe and cease to provide liquidity on that venue. If costs of entry in liquidity provision are sufficiently low that the marginal liquidity provider makes close to zero expected profit, then the marginal provider will rationally withdraw from that particular venue when fees are introduced.

For informed traders who seek to earn positive expected profits by trading on private information, the decision to subscribe to data fees has some similar features to the decision for market makers. The expected benefit of subscribing for informed traders can come from enhanced information about the state of the market-wide order book, which can be informative about the value of the asset (e.g. if limit orders contain information about fundamental value as in Brogaard, Hendershott, and Riordan 2019). The benefit of subscribing can also come from the ability to condition order routing decisions based on the state of the order book and design trading strategies that minimize costs accordingly. The fee-introducing venues that we study have sizeable market shares (around 10% on average), but it is not obvious that informed traders need to condition on the limit orders placed on these venues to better estimate the fundamental value of the security. Clearly trading costs are a very important consideration for informed traders, as these costs act as a tax on the expected return from information acquisition. If having order book information can help minimize

these costs, then informed traders may choose to pay the fees. The fees themselves however also act as a (lump sum) tax on expected profits, so informed traders may well choose to not pay this tax and route orders to competitor venues after fees are introduced. A break-even condition when information acquisition is costly would imply that the marginal informed trader will not be able to pay the fee and still earn non-negative profits.

In many standard market microstructure models such as Kyle (1985), liquidity traders are assumed to either demand or supply shares in a way that is inelastic to price. In other words, these traders submit a market order regardless of price or trading costs. While this is clearly a highly stylized model of uninformed trader behavior, it is useful to consider how fees for order book information may affect such a trader. Presumably, such a liquidity trader with inelastic demand or supply would not choose to pay a lump sum tax in the form of order book data fees when they choose not to condition their orders on current information. They may still route orders some fraction of their orders to the venue (either directly or via an intermarket sweep order) or may exclude the venue from their routing decisions. A liquidity trader could instead rationally choose to subscribe if they believe doing so will allow them to trade cheaply on the venue and this leads to lower overall transaction costs compared with not having this information. The rational liquidity trader would have to analyze the costs and benefits of the order book data in a similar way to the informed trader and decide accordingly.

We do not take a prior position regarding which (if any) of these trader types may be more sensitive than the other. Instead we try to map how three variables (average trading costs, the informativeness of trade arrivals and market maker profits from liquidity provision) are likely to change if one group of traders has a higher propensity to not to subscribe.

First, for informed traders, whose trades contribute to the price discovery process, we

argue that if they are less likely to subscribe to order book fees and reroute their orders elsewhere, then the average price impact per trade is expected to fall for venues introducing fees. Market makers subsequently face less adverse selection when providing liquidity and earn higher revenue per trade from market making, all else being equal. Average effective spreads may decrease or remain unchanged depending on the degree of competition between market makers.

Analogous but opposite arguments apply to liquidity traders, whose trades by definition carry no informational content. If this category is less likely to subscribe to order book fees and therefore route their orders elsewhere, then average price impact per trade is expected to increase for venues introducing fees. Market makers are more likely to trade against informed traders and therefore face more adverse selection when providing liquidity than prior to fees. Average effective spreads must increase to compensate for the higher costs of providing liquidity, all else being equal. Revenue per trade either falls or is unchanged depending on the degree of competition between market makers.

Lastly, if market makers are relatively more sensitive to the cost of order book information, then the degree of competition in the provision of liquidity is expected to fall as these market makers can no longer compete on a level playing field with those that do subscribe. Average market maker revenue per trade (i.e. for the market makers that do subscribe) and average trading costs increase after fees are introduced. Price impact may increase or decrease depending on whether informed or liquidity traders are more sensitive to trading costs. If informed traders are more (less) sensitive to trading costs, then price impact is expected to fall (rise).

We summarize these three mechanisms and their predictions for trading costs, market making profits and trade price impact in Table 1.

Table 1 about here

4 Data and Summary Statistics

Data for our analysis is obtained from the NYSE Trade and Quote (TAQ) database. The TAQ data contain trades and quotes for all securities listed on the NYSE, the American Stock Exchange (AMEX), the Nasdaq National Market System and Small Cap issues. Using the TAQ data, we calculate a set of liquidity and trading activity variables at the stock-day level separately for the ATS introducing the fee and for all other trading venues in three month windows either side of the fee introduction date. The sample includes all stocks in the S&P 500, S&P 400 and S&P 600 indices (collectively, the S&P 1500 index). Ticker identifiers for S&P 1500 members over the relevant sample periods are obtained from Thomson Reuters Tick History, maintained by Refinitiv.

The set of variables we analyse are the log dollar volume traded by stock-day-venue and the average percent effective spread, percent realized spread and percent price impact across trades by stock-day-venue. Log dollar volume is defined as the log of the total dollar value traded across the N_{ijt} trades in stock i on venue j and date t :

$$\text{Log Dollar Volume}_{ijt} = \ln \sum_{n=1}^{N_{ijt}} \text{vol}_{ijt}^n \quad (1)$$

The effective spread, realized spread and price impact for trade n in stock i on venue j and

date t are respectively defined as:

$$ES_{ijt}^n = 100 \times \frac{D_{ijt}^n (P_{ijt}^n - Mid_{ijt}^n)}{Mid_{ijt}^n} \quad (2)$$

$$RS_{ijt}^n = 100 \times \frac{D_{ijt}^n (P_{ijt}^n - Mid_{ijt}^{n+5})}{Mid_{ijt}^n} \quad (3)$$

$$PI_{ijt}^n = 100 \times \frac{D_{ijt}^n (Mid_{ijt}^{n+5} - Mid_{ijt}^n)}{Mid_{ijt}^n} \quad (4)$$

where D_{ijt}^n is an indicator variable equal to one if the n^{th} trade in stock i on venue j and date t is a customer buy and negative one if the trade is a customer sell, P_{ijt}^n is the trade price, Mid_{ijt}^n is the prevailing midquote at the time of the trade and Mid_{ijt}^{n+5} is the prevailing midquote five minutes after the trade time. Customer trade directions are assigned according to the Lee and Ready (1991) algorithm. For each trade-level variable defined in Equations (2), (3) and (4), our stock-day-venue value is the average across all trades in that stock-day-venue.³ All variables are Winsorised at the 1% level to limit the effect of outliers on our results.

To illustrate how we define our variables, ARCA first introduced fees for order book information on January 1, 2009. For every stock in the S&P 1500 index and every trading date in a three month window either side of this date, we calculate each activity or liquidity variable for all trades and quotes on ARCA and also for all trades and quotes on all other trading venues excluding ARCA. We then use the stock-day averages of trading activity and liquidity variables for trades and quotes on ARCA as our treated group and stock-day averages using all trades and quotes from other venues as the control group. We do this for each of our three events corresponding to fee introduction on ARCA, Direct Edge and BATS respectively. Our final sample consists of 6,198 stock-venue combinations corresponding to

³We gratefully use SAS code provided by Craig Holden and Stacey Jacobsen to calculate each of our variables, as described in Holden and Jacobsen (2014), with slight modifications to separate estimation into different venues.

trades in S&P 1500 stocks on ARCA, BATS and Direct Edge and the control venues, across 368 trading days around each fee introduction event.⁴ Table 2 contains summary statistics for each variable across the entire pooled sample and averages split by before and after fee introduction.

Table 2 about here

Average volume traded across all stock-day-venues rises slightly from around \$32.6m before fee introduction to \$35.3m in the post-fee introduction period. Quoted spreads trend slightly lower from around 36 bps in the pre-event period to 32.5bps in the post-event period, while dollar depth are largely unchanged. Average trade costs (effective spreads) are very similar in the pre- and post-event periods, though there are small decreases in average realized spreads and increases in price impact. These summary statistics primarily capture general trends in trading activity across the three sample periods.

Table 3 presents summary statistics split by venue type: ATS or other. Panel A contains averages across stock-days for trading on ATS while Panel B contains averages across stock-days on other venues. Trading activity on the ATS introducing the fees represents approximately 12% of total volume traded across all stock-venue-days in our sample with \$8.3m worth of stock traded on the ATS on average compared with \$59.6m on other venues. There is also less liquidity at the quotes for the ATS compared with across all other venues (as measured by total depth available at the best bid and offer and the quoted spread at the best bid and offer). The average quoted spread across all stock-day-venues is 34 bps, however the corresponding levels are 48bps for the ATS and 20 bps for other venues. The average depth at the best bid and ask is \$17,600 across all stock-venue-days, but is \$8,890

⁴The final number of stock venues is slightly larger than 1500×4 because some members of the index have dual listings.

and \$23,700 for the ATS and the other venues respectively. The control group aggregates total liquidity across multiple venues (including the primary exchange for that stock and all other trading venues). It is not surprising that liquidity as measured at the quotes appears to be worse on the ATS than in the control group but due to aggregation across venues for the control group, this does not indicate that the ATS has substantially poorer liquidity at the quotes than any other specific venue.

Table 3 about here

There are generally small differences in average effective spreads, realized spreads or price impact between the ATS (treated) and other venues (control). The average effective spread is close to 11bps for both the treated and control, which on average can be decomposed into 7bps due to price impact and 4bps that accrues to the market maker as realized spread. Effective spreads trend slightly higher for trades on both the ATS and the control venues, though the rise is larger for control venues. Realized spreads increase and price impact falls on both venue categories, but the effects are more pronounced for the ATS relative to the control variables. Prior to fee introduction, average realized spreads are approximately 0.6bps lower on the ATS while after the fee introduction, they are 0.3bps smaller. Price impact is on average 0.2bps higher for trades on the ATS prior to fee introduction but are 0.2bps.

Figure 1 plots the average trading activity in the pre- and post-event periods. The y-axis depicts the average of log dollar value traded by venue type while the x-axis depicts the number of days before and after the fee-introduction date. The “treated” line is the daily simple average of log dollar volume on the ATS introducing the fee across all stocks in our sample. The “control” line is the equivalent average across the same stocks, but traded on all other venues.

In-line with Tables 2 and 3, average volumes traded are significantly lower on the ATS in both the pre-event and post-event periods, and also appear to be trending slightly lower in both categories over the pre-event sample period, reflecting a general trend in total trading activity that is unlikely to be related to the introduction of fees.

Figure 1 about here

Figure 2 depicts the average effective spreads by venue type in the pre- and post-event periods. Before fees are introduced, there is no persistent notable gap between the two series. After fees are introduced, we observe a small but persistent gap between the average trading cost for a trade on ATS introducing the fees and the control group. The two series track each other very closely over the pre- and post-event periods.

Figure 2 about here

Plots of the effective spread decomposition are contained in Figure 3. Panel A contains market maker realized spreads while Panel B contains price impact. Realized spreads are persistently lower for trades on the ATS prior to fee introduction and this gap falls after fees are introduced. Price impacts are slightly higher for trades on the ATS before the fee introduction and slightly lower after. For all four variables, the time-series of stock-day averages across venues are highly correlated and appear to be following common trends.

Figure 3 about here

5 Do Market Participants Respond to Data Fees?

We use a difference-in-differences approach to causally identify the effect of fees for order book information on trading volumes and trading costs. Our strategy treats the introduction of fees as natural experiments that exogenously raise the price for order book information on one exchange but not others. For each of our three events, we construct a sample of stock-day observations of trading volume and costs for the venue introducing the fee and a sample of stock-day observations for all other venues where the stock is traded. We assign stock-days on the venue introducing fees as the treated group and stock-days on all other venues as the control group. We then pool the sample together across all three events and estimate the effect of fee introduction using both a standard difference-in-differences approach and also using two-way fixed effects defined at the stock-venue and day level. The standard difference-in-differences regression is given by:

$$y_{ijt} = \alpha + \beta_1 post_t + \beta_2 D_j + \beta_3 D_j \times post_t + \varepsilon_{ijt} \quad (5)$$

where y_{ijt} is the trading activity or market quality for stock i on venue j (either the ATS introducing the fees or the sum of all other venues) and date t , $post_t$ is a dummy taking the value 1 if the date is after the fee introduction for venue j , D_j is a dummy variable taking the value 1 if the venue j is the venue introducing the fees. Equation (5) is the standard difference-in-differences estimator, as in Meyer (1995) amongst others. The parameter β_3 identifies the causal effect of fee introduction in the absence of confounding trends or selection to treatment status based on unit-specific unobservables. In our case units are determined at the stock-venue level and so selection into the treatment or control group cannot plausibly depend on unobservable effects. However, since we have daily panel data we also estimate a

model with two-way stock-venue and date fixed effects. This model is given by:

$$y_{ijt} = \alpha_{ij} + \gamma_t + \beta \times D_{jt} + \varepsilon_{ijt} \quad (6)$$

where α_{ij} is a stock-venue fixed effect, γ_t is a daily fixed effect, D_{jt} is a dummy variable taking the value one if the venue is the ATS and the trading date is equal to or later than the fee introduction date and ε_{ijt} is an error term. The model described in Equation (6) nests the simpler expression in Equation (5) while allowing for more flexible time effects and unobserved components at the stock-venue level. The key parameter of interest is β which again captures the causal effect of the treatment (fee introduction) on the trading activity and market quality for trading on the ATS relative to the control group.

In both models, we pool our data across all three events and estimate a single treatment effect corresponding to their average effect of fee introduction on trading activity and market quality. Pooling across the three events helps to limit the reliance on parallel trends insofar as similar confounding trends would need to be present across all three events to contaminate our results in a systematic way. The evidence presented in Figures 1, 2 and 3 strongly suggests that trading volumes, trading costs and trading cost components closely track each other over both the pre-event and post-event period, supporting the parallel trends identifying assumption. We estimate both models using standard heteroskedasticity-robust standard error estimators and two-way cluster robust estimators where clustering is defined at the stock and date level.

Estimates from our differences-in-differences regressions for log dollar value traded are presented in Table 4. These regressions test whether venues that introduce fees for order book information suffer a fall in market share relative to other venues. Our empirical approach

is designed to causally identify the effect of fees on market share and execution quality. As we argue in Section 3 a change in market share reflects that at least some traders do not subscribe to the fees and prefer to route their orders to other exchanges, relative to the pre-fee environment. Our log volume regressions are therefore tests of the hypothesis that some traders are sensitive to the cost of order book information and exchanges face downward sloping demand curves for order book information.

Column (1) refers to the simple difference-in-differences specification corresponding to Equation (5) with heteroskedasticity robust standard errors. Column (2) contains equivalent estimates but using cluster robust standard errors with clustering on stock and date. Column (3) contains the estimates from our panel difference-in-differences specification (i.e. with stock-venue and date fixed effects) corresponding to Equation (6) and heteroskedasticity robust standard errors. Column (4) is the same specification as Column (3) but with two-way cluster robust standard errors.

Table 4 about here

The point estimates for the effect of fee introduction on log volume in Table 4 are highly consistent across specifications — venues lose 5-8% of their existing volume on average after introducing fees, relative to venues that do not introduce fees.⁵ As these regressions are simple difference-in-differences, the point estimates in Columns (1) and (2) are almost identical to the difference-in-differences in averages across venues and pre-post windows in Table 3. There are very small discrepancies due to an uneven number of trading days in the pre- and post windows due to weekends and public holidays. The estimates in Columns (3) and (4) allow for stock and time fixed effects and so nest the specifications in Columns (1) and

⁵Since the dependent variable is in logs, the treatment effect parameter captures the difference-in-differences in logs and so captures percentage changes in volumes.

(2). The effect is highly significant under all specifications other than the simple difference-in-differences with clustered standard errors. Column (4) is the most flexible specification and has standard errors that are robust to arbitrary correlation within stock and time dimensions. For these reasons it is our most preferred specification and importantly the effect is highly significant for this regression. In terms of dollar value, a 5-8% fall corresponds to around \$440,000 to \$700,000 less daily ATS volume for the average stock in our sample, or \$80,000 to \$120,000 for the median stock. The coefficient estimates in Table 4 demonstrate that at least some market participants are sensitive to the introduction of fees and prefer to route their orders to other venues after fees are introduced.

6 Which Market Participants Respond to Data Fees?

Given that venues face a downward sloping demand curve for order book information, we next ask whether some groups of traders are more sensitive to fees than others. We focus on three categories of traders that are present in many canonical market microstructure models: informed traders, liquidity traders and market makers. To distinguish between these different categories, we examine how the average informativeness of customer trades (price impact), dealer revenue from market making (realized spreads) and total trading costs (effective spreads) change around fee introduction. Section 3 discuss how these variables are predicted to respond if one category of traders is more sensitive to fees than others.

Due to the complex interactions between informed trading, trading costs and venue selection, no single variable in isolation can perfectly distinguish between one trader category over the other two. We choose to first examine average price impact because detecting a significant positive or negative effect of fees on this variable can most clearly rule out either the informed or liquidity trader categories. We argue that observing a significant increase in

the average price impact of trades on the venue introducing fees cannot be consistent with informed traders being the most sensitive to data fees. If informed traders are more sensitive to fees than liquidity traders and therefore route fewer orders to the venue introducing the fee, average price impact of trades should fall. If however liquidity traders are more sensitive to fees and reroute their orders to other venues, the average trade on the venue introducing fees should have higher price impact. If our third category of traders (market makers) are more sensitive to fees, we expect dealer revenues to increase and effective spreads to (weakly) increase but price impact could either rise or fall depending on whether liquidity or informed traders are more sensitive to trading costs. Significant effects on price impact can therefore rule out either informed or liquidity traders as the most sensitive to fees, but alone cannot rule-out the market maker trader category.

Parameter estimates for the effect of fees on average price impact are presented in Table 5. In each of the four specifications, the effect of fees on average price impact is negative and highly significant. The size of the parameter estimate is approximately -0.4 bps across all four specifications, which corresponds to approximately 6% of the average trade price impact across all stock-day-venues in our sample. The size of the *t*-statistic is well in excess of the 1% threshold in all four specifications.

Table 5 about here

Table 5 demonstrates that trades on ATS become less informed on average after the introduction of data fees, relative to trades on other venues with no change in data fees. The evidence on price impact suggests that liquidity traders routing decisions are not more sensitive to order book fees than other trader categories. Liquidity traders either are more willing to pay for order book information than market maker or informed traders, or are less

reliant on order book information in deciding where to route their orders and so continue to use the ATS that introduce fees despite not observing the state of the ATS order book.

Table 5 cannot directly distinguish between market makers or informed traders being more sensitive to data fees. Arguably, the simplest explanation for decreased price impact is that informed traders route their orders to other venues after fees are introduced. In this case, market makers face less adverse selection costs when providing liquidity. Depending on the degree of competition in liquidity provision, market makers will pass this saving onto customers via lower effective spreads and / or earn higher revenue per trade due to lower their lower costs.

Decreased price impact is also consistent with an alternative explanation based on the behavior of market makers. If some market makers choose not to subscribe to order book information, they will suffer from information asymmetry with market makers who do subscribe. Standard Winner's Curse arguments can induce these market makers to compete less aggressively when providing liquidity on the venue introducing fees which leads to higher trading costs and dealer revenue per trade. Price impact could plausibly increase in response to higher trading costs if informed traders are more sensitive to these costs than liquidity traders.

Whether informed traders or market makers are relatively more sensitive to order book fees is therefore distinguished by the impact of data fees on effective spreads, which measure trading costs, and realized spreads, which measure average dealer revenue per trade. If informed traders are more likely to route their orders elsewhere after fees, then either effective spreads should decrease, realized spreads should increase or both. If market makers are choosing not to subscribe, thus lowering the degree of competition between market makers, effective spreads and realized spreads should increase. If informed traders are also more

sensitive to trading costs than uninformed traders, then price impact will fall, consistent with Table 5.

Table 6 contains parameter estimates for the effect of data fees on effective spreads and realized spreads. Coefficient estimates for effective spreads show a very small reduction in average trading costs on the venues introducing fees for order book information relative to other venues. The economic magnitude of the effect is around 0.05 bps per trade, compared with a pre-fee average of around 11bps. Though the statistical significance of the effect is mixed, in our most flexible specifications, we can reject the null hypothesis that the effective spread increases after fees at the 5% level or better. For realized spreads, we detect a strongly significant positive effect of approximately 0.3 bps per trade. The size of the effect is highly consistent and the magnitude of the *t*-statistics exceed the 1% thresholds across specifications.

Table 6 about here

In Table 6, the reduction in the costs of market making due to reduced average price impact almost entirely accrue to the market makers in the form of higher realized spreads, with only a very small fraction of the savings (if any) being passed on to customers. These patterns combined with the fall in total volume (Table 4) and average price impact (Table 5) are consistent with the trading decisions of informed traders being more sensitive to data fees than liquidity traders or market makers. Informed traders appear to be less likely to pay the fixed costs of subscribing to data fees and then choose to route their order to other venues after fees are introduced. Market makers capture higher profits on a per-trade basis, though their total revenue from market making on the venue introducing the fee is also affected by lower volumes. Crucially, Table 6 combined with the fall in average price impact in Table

5 are inconsistent with the key effect of fees being less market maker competition. In this scenario, a fall in price impact requires first for average trading costs to rise and for this rise to disproportionately affect order routing decisions of informed traders. Table 6 indicates that effective spreads are either lower or unchanged after fees are introduced. Even if routing decisions of informed traders are more sensitive to trading costs, we find no evidence that these increased on the venue introducing fees relative to the pre-fee environment.

7 Robustness

We have argued that the effect of fees on trading volume and trade costs is consistent with fees having a disproportionate effect on the order placement of informed traders compared with liquidity traders or market makers. Ideally, trade-level data that identified individual traders and whether they trade for informational reasons vs. liquidity reasons would help verify our conclusions.

As we do not have such a dataset, we instead try to provide corroborating evidence by considering the effect of data fees for stocks where informed trading is relatively more or less important. If we are able to identify a group of stocks where informed trading makes up a relatively greater proportion of overall trading, then under our interpretation of Tables 4 — 6, we expect these stocks to have greater reductions in volume and price impact and a greater increase in realized spreads after fees are introduced. In the absence of indicators for informed vs. liquidity trading, we rely on splitting our sample of stocks into two sub-samples based on average price impact across all venues in the pre-fee period. Stocks with higher than median price impact are grouped into the “High Price Impact” category and those with below median price impact are grouped into the “Low Price Impact” category.

If informed traders disproportionately reduce their volume relative to other traders, then

we expect to see larger reductions in trading volume and price impact combined with greater realized spreads in the “High Price Impact” stocks relative to the other group, as these traders are assumed to constitute a greater fraction of trading in the High category. Table 7 contains parameter estimates for all four dependent variables split by average pre-fee price impact. Panel A contains the results for high price impact stocks. Panel B contains the results for low price impact stocks.

Table 7 about here

Stocks in the high price impact category (Panel A) experience greater falls in volume, effective spreads and price impact than stocks in the low price impact category (Panel B). Stocks in the high price impact category also experience greater increases in realized spreads than stocks in the low price impact category. For volume traded, the estimated effect in high price impact stocks corresponds to between an 8% and 10% reduction, while price impact falls by 0.6bps and realized spreads rise by 0.5bps. For low price impact stocks, the corresponding numbers are 3% and 6% for volume traded and 0.2bps for price impact and realized spreads.

Under the assumption that stocks with higher average price impact have higher average informed trading than stocks with low price impact, these results provide further support to the informed trader mechanism. The fact that the signs on the coefficients are consistent across sub-samples and regression specifications is also important in our context. Fees are charged at a market-wide level. Traders cannot subscribe to orderbook information for some stocks and not for others. Therefore, if some informed traders choose not to subscribe, then they face similar incentives regarding routing their orders in the post-fee environment across all stocks, as exhibited in Table 7.

In addition to our main specifications, we also estimate the effect of fee introduction for each venue and dependent variable separately and estimate our pooled difference-in-differences regressions in one month windows around each event rather than the three month window. Table 8 contains the difference-in-differences estimates for regressions estimated separately by exchange. Column (1) of Table 8 contains the treatment effect estimate for the introduction of fees on ARCA. The treated group in Column (1) are constructed from trades on ARCA while the control group contains trades on all other venues. Columns (2) and (3) are defined analogously but for the introduction of fees on Direct Edge and BATS respectively. We only report the effects for the most general specifications corresponding to Equation (6) and with two-way cluster robust standard errors.

Table 8 about here

The signs of parameters in Table 8 are consistent across all three venues, suggesting that the direction of the effects of fees is consistent on all three venues. Log volume undergoes a large and significant fall for two of the three exchanges (ARCA and BATS). For Direct Edge, the point estimate is also negative but is smaller and is statistically insignificant. Price impact falls by a statistically significant amount on all three venues, though the effect is largest on ARCA. The point estimate for effective spreads is negative on all three venues but is generally small and statistical significance is mixed. The effect of fees on realized spreads are positive for all three venues but only statistically significant for ARCA and BATS.

Finally, our results are consistent when we use a one month window for all variables other than for log volume. In our most robust panel specification, we detect a much smaller and insignificant effect on volume traded. Under the standard difference-in-differences specification,

the coefficient is positive but also insignificant. For trading costs and the decompositions variables, the estimates are similar to our main effects. There is a consistent increase in dealer realized spreads and falls in effective spreads and price impact, though the magnitude of the effect on price impact and realized spreads are smaller in magnitude than our main results. It is possible that it takes some time for traders to adjust to the new equilibrium of wider spreads and lower market maker competition on the ATS introducing the fees, which could explain the insignificant effect over the one month pre-post event windows.

Table 9 about here

8 Conclusion

This paper explores the role of data fees on exchange competition. While exchange data has properties of a public good, the exchanges control it and can, with the approval of the Securities and Exchange Commission, set a price on it. We show that when a trading venue introduces a fee on its data, it experiences a decline in trading activity. This shows that the demand for such data is not perfectly inelastic and so exchanges experience at least some trade off between higher volume and higher data fees.

By examining different measures of liquidity we can infer the type of trader that is predominantly impacted by the data fee introduction. We find that the introduction of a data fee results in price impact falling and realized spread increasing, with negative but small changes to effective spreads. These changes are most consistent with informed traders decreasing their trading on the venue increasing a data fee. We speculate that this change does not materially change the amount of information acquired, simply the distribution of where the informed traders choose to trade on the information.

References

- Alan, N. S. and R. A. Schwartz (2013). Price discovery: The economic function of a stock exchange. *Journal of Portfolio Management* 40(1), 124–132.
- Brogaard, J., T. Hendershott, and R. Riordan (2019). Price discovery without trading: Evidence from limit orders. *The Journal of Finance* 74(4), 1621–1658.
- CBOE (2012). Notice of Filing and Immediate Effectiveness of Proposed Rule Change Relating to Fees Associated with Receipt of the EDGX Book Feed. https://cdn.cboe.com/resources/regulation/rule_filings/approved/2012/SR-EDGX-2012-14-Revised.pdf. Accessed: 2020-09-22.
- Hendershott, T. and C. M. Jones (2005). Island goes dark: Transparency, fragmentation, and regulation. *The Review of Financial Studies* 18(3), 743–793.
- Hendershott, T., M. Rysman, and R. Schwabe (2020). Stock exchanges as platforms for data and trading. *Working Paper*.
- Holden, C. W. and S. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *The Journal of Finance* 69(4), 1747–1785.
- Jones, C. M. (2018). Understanding the market for u.s. equity market data. *Working Paper*.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 1315–1335.
- Lee, C. M. and M. J. Ready (1991). Inferring trade direction from intraday data. *The Journal of Finance* 46(2), 733–746.
- Meyer, B. D. (1995). Natural and quasi-experiments in economics. *Journal of business & economic statistics* 13(2), 151–161.
- MFA and AIMA (2018). Petition for rulemaking regarding market data fees and request for guidance on market data licensing practices.
- Pagnotta, E. S. and T. Philippon (2018). Competing on speed. *Econometrica* 86(3), 1067–1115.
- Ramos, S. B. and E.-L. von Thadden (2008). Stock exchange competition in a simple model of capital market equilibrium. *Journal of Financial Markets* 11(3), 284–307.
- Reuters (2013). BATS exchanges to start charging for market data. <https://www.reuters.com/article/batsglobalmarkets-marketdata-fees-idUSL2NOD52AC20130418>. Accessed: 2020-09-22.

SEC (2008). Self-Regulatory Organizations; NYSE Arca, Inc.; Order Setting Aside Action by Delegated Authority and Approving Proposed Rule Change Relating to NYSE Arca Data. <https://www.sec.gov/rules/sro/nysearca/2008/34-59039.pdf>. Accessed: 2020-09-22.

Figure 1: ATS Market Share Around Fee Introduction

This figure plots the average log dollar volume traded across the stocks in our sample in three month windows around fee introduction, split by venues introducing fees and other venues. The y-axis reports the average log dollar value traded across all S&P1500 stocks by day, pooled across the three fee introduction events. The x-axis reports the number of trading days before and after the introduction of fee. Log dollar volume by day traded on the three ATS introducing the fee is depicted in the “treated” line. Log dollar volume by day traded on all other venues not introducing fees is depicted in the “control” line.




Figure 2: Effective Spreads Around Fee Introduction

This figure plots the average effective spreads across the stocks in our sample in three month windows around fee introduction, split by venues introducing fees and other venues. The y-axis reports the average effective spread for trades in all S&P1500 stocks by day, pooled across the three fee introduction events. The x-axis reports the number of trading days before and after the introduction of fee. Effective spreads on the three ATS introducing the fee is depicted in the “treated” line. Effective spreads on all other venues not introducing fees is depicted in the “control” line.





Figure 3: Spread Decomposition Around Fee Introduction

This figure plots the average realized spreads (Panel a) and price impact (Panel b) across the stocks in our sample in three month windows around fee introduction, split by venues introducing fees and other venues. The y-axis reports the average realized spread and price impact for trades in all S&P1500 stocks by day, pooled across the three fee introduction events. The x-axis reports the number of trading days before and after the introduction of fee. Realized spreads and price impact on the three ATS introducing the fee are depicted in the “treated” line. Realized spreads and price impact on all other venues not introducing fees are depicted in the “control” line.

(a) Realized Spreads



(b) Price Impact




Table 1: Predicted spread decompositions vs. trader categories

This table contains predictions for changes in price impact, effective spreads and realized spreads when different trader categories are assumed to be relatively more sensitive to the introduction of order book fees than the other categories. The three trader categories are market makers, informed traders and liquidity traders.

Category	Price Impact	Effective Spreads	Realized Spreads
Market Makers	↑ / ↓	↑	↑
Informed Traders	↓	↑ / ↓	↑
Liquidity Traders	↑	↑	↑ / ↓

Table 2: Stock-Day Trading Summary Statistics — All Venues

This table reports means, standard deviations, minimums, maximums and 25th, 50th and 75th quantiles for trading activity and liquidity variables for our sample, split into trading on ATS and all other venues. The sample includes all stock-venue-day combinations for S&P 1500 members over three month windows around the introduction of fees on ARCA, BATS and Direct Edge, where venues are divided into the ATS introducing the fee and all other trading venues. Dollar volume traded is the total volume traded by stock-day-venue measured in millions of dollars. The quoted spread is the time-weighted average of the difference between the best bid and ask, divided by the midquote (expressed in %). Dollar depth is the time-weighted average of the dollar volume available at the best bid and ask measured in thousands of dollars. Effective spread, realized spread and price impact are defined as per Section 4. All variables are winsorized at the 1% level.

	Mean	SD	Min	25%	50%	75%	Max	Mean Pre-Fee	Mean Post-Fee
Dollar Value Traded (mil)	33.932	87.523	0.0158	1.1414	5.5106	24.318	747.66	32.612	35.282
Log Dollar Value Traded	15.440	2.1801	9.6649	13.947	15.522	17.006	20.432	15.370	15.511
Quoted Spread (%)	0.3430	0.4433	0.0205	0.0882	0.1890	0.4099	2.9867	0.3601	0.3254
Dollar Depth (000)	17.652	31.091	0.6608	5.1290	9.0931	16.790	275.76	17.730	17.572
Log Dollar Depth	9.1751	1.0073	6.4934	8.5427	9.1153	9.7285	12.527	9.1798	9.1702
Effective Spread (%)	0.1088	0.1332	0.0134	0.0385	0.0684	0.1252	1.0751	0.1091	0.1085
Realized Spread (%)	0.0354	0.0849	-0.2399	0.0050	0.0205	0.0485	0.6508	0.0378	0.0330
Price Impact (%)	0.0701	0.1069	-0.1048	0.0152	0.0380	0.0857	0.7734	0.0678	0.0725

Table 3: Stock-Day Trading Summary Statistics — Split by ATS and Other Venues

This table reports means, standard deviations, minimums, maximums and 25th, 50th and 75th quantiles for trading activity and liquidity variables for our sample. The sample includes all stock-venue-day combinations for S&P 1500 members over three month windows around the introduction of fees on ARCA, BATS and Direct Edge, where venues are divided into the ATS introducing the fee and all other trading venues. Dollar volume traded is the total volume traded by stock-day-venue measured in millions of dollars. The quoted spread is the time-weighted average of the difference between the best bid and ask, divided by the midquote (expressed in %). Dollar depth is the time-weighted average of the dollar volume available at the best bid and ask measured in thousands of dollars. Effective spread, realized spread and price impact are defined as per Section 4. All variables are winsorized at the 1% level.

	Mean	SD	Min	25%	50%	75%	Max	Mean Pre-Fee	Mean Post-Fee
Panel A: Trading on ATS									
Dollar Value Traded (mil)	8.2859	29.215	0.0158	0.4045	1.5831	6.1810	747.66	8.9577	7.6291
Log Dollar Value Traded	14.252	1.9151	9.6649	12.910	14.274	15.637	20.432	14.338	14.168
Quoted Spread (%)	0.4830	0.5396	0.0205	0.1311	0.3049	0.6070	2.9867	0.4484	0.5168
Dollar Depth (000)	11.524	18.954	0.6608	4.1652	7.1461	12.153	275.76	11.340	11.705
Log Dollar Depth	8.8868	0.8940	6.4934	8.3345	8.8743	9.4053	12.527	8.8737	8.8996
Effective Spread (%)	0.1069	0.1337	0.0134	0.0367	0.0655	0.1225	1.0751	0.1067	0.1070
Realized Spread (%)	0.0330	0.0937	-0.2399	-0.0019	0.0192	0.0520	0.6508	0.0298	0.0362
Price Impact (%)	0.0701	0.1179	-0.1048	0.0092	0.0349	0.0894	0.7734	0.0735	0.0668
Panel B: Trading on Other Venues									
Dollar Value Traded (mil)	59.578	114.68	0.0158	4.9669	17.268	60.206	747.66	61.607	57.594
Log Dollar Value Traded	16.628	1.7368	9.6649	15.418	16.664	17.913	20.432	16.684	16.573
Quoted Spread (%)	0.2030	0.2501	0.0205	0.0673	0.1305	0.2475	2.9867	0.2025	0.2035
Dollar Depth (000)	23.779	38.717	0.6608	6.5937	11.781	23.112	275.76	23.804	23.754
Log Dollar Depth	9.4634	1.0314	6.4934	8.7939	9.3743	10.048	12.527	9.4668	9.4601
Effective Spread (%)	0.1107	0.1327	0.0134	0.0406	0.0714	0.1278	1.0751	0.1103	0.1111
Realized Spread (%)	0.0378	0.0750	-0.2399	0.0090	0.0214	0.0456	0.6508	0.0362	0.0394
Price Impact (%)	0.0701	0.0947	-0.1048	0.0194	0.0404	0.0826	0.7734	0.0714	0.0688

Table 4: Difference-in-Differences Regressions: Volume

This table reports coefficients (t -statistics) for difference-in-differences regressions of log dollar volume around fee introduction. The treated sample includes daily trading in all S&P1500 stocks on the venue introducing order book fees in three month windows around fee introduction. The control sample includes the daily trading in the same stocks on all other venues not introducing fees. Each stock-day-venue combination contributes one observation to the regression. Observations across each of the three fee introduction events are pooled into a single regression. Column (1) is a simple differences-in-differences specification with a common intercept, single post-period time dummy and heteroskedasticity robust standard errors. Column (2) is the same specification but with two-way clustered standard errors clustered at the stock-day level. Column (3) includes separate stock-venue and day fixed effects and uses heteroskedasticity robust standard errors. Column (4) is the same as Column (3) but with two-way clustered standard errors.

	(1)	(2)	(3)	(4)
	Model A	Model B	Model C	Model D
Log Dollar Value Traded	-0.0583 (-8.27)	-0.0583 (-1.07)	-0.0742 (-33.6)	-0.0742 (-8.05)
N	6,198	6,198	6,198	6,198
T	368	368	368	368
Stock-Venue FE	-	-	X	X
Time FE	-	-	X	X
Error Type	HC1	Clust.	HC1	Clust.

Table 5: Difference-in-Differences Regressions: Price Impact

This table reports coefficients (*t*-statistics) for difference-in-differences regressions of average price impact around fee introduction. The treated sample includes trades in all S&P1500 stocks on the venue introducing order book fees in three month windows around fee introduction. The control sample includes the trades in the same stocks on all other venues not introducing fees. Each stock-day-venue combination contributes one observation to the regression. Observations across each of the three fee introduction events are pooled into a single regression. Column (1) is a simple differences-in-differences specification with a common intercept, a single post-period time dummy and heteroskedasticity robust standard errors. Column (2) is the same specification but with two-way clustered standard errors clustered at the stock-day level. Column (3) includes separate stock-venue and day fixed effects and uses heteroskedasticity robust standard errors. Column (4) is the same as Column (3) but with two-way clustered standard errors.

	(1)	(2)	(3)	(4)
	Model A	Model B	Model C	Model D
Price Impact (%)	-0.0040	-0.0040	-0.0042	-0.0042
	(-9.79)	(-4.33)	(-14.6)	(-8.04)
<i>N</i>	6,198	6,198	6,198	6,198
<i>T</i>	368	368	368	368
Stock-Venue FE	-	-	X	X
Time FE	-	-	X	X
Error Type	HC1	Clust.	HC1	Clust.

Table 6: Difference-in-Differences Regressions: Realized & Effective Spreads

This table reports coefficients (t -statistics) for difference-in-differences regressions of average realized and effective spreads around fee introduction. The treated sample includes trades in all S&P1500 stocks on the venue introducing order book fees in three month windows around fee introduction. The control sample includes the trades in the same stocks on all other venues not introducing fees. Each stock-day-venue combination contributes one observation to the regression. Observations across each of the three fee introduction events are pooled into a single regression. Column (1) is a simple differences-in-differences specification with a common intercept, a single post-period time dummy and heteroskedasticity robust standard errors. Column (2) is the same specification but with two-way clustered standard errors clustered at the stock-day level. Column (3) includes separate stock-venue and day fixed effects and uses heteroskedasticity robust standard errors. Column (4) is the same as Column (3) but with two-way clustered standard errors.

	(1)	(2)	(3)	(4)
	Model A	Model B	Model C	Model D
Effective Spread (%)	-0.0005 (-1.00)	-0.0005 (-1.04)	-0.0006 (-2.86)	-0.0006 (-2.56)
Realized Spread (%)	0.0032 (9.71)	0.0032 (3.46)	0.0032 (11.5)	0.0032 (5.30)
N	6,198	6,198	6,198	6,198
T	368	368	368	368
Stock-Venue FE	-	-	X	X
Time FE	-	-	X	X
Error Type	HC1	Clust.	HC1	Clust.

Table 7: Difference-in-Differences Regressions: Price Impact Categories

This table reports coefficients (t -statistics) for difference-in-differences regressions for log dollar volume traded, effective spread, realized spread and price impact around fee introduction, split by average pre-fee price impact. All details are identical to Tables 4, 5 and 6 except that we split the sample into two categories of stocks based on the average price impact of trades in the pre-fee environment. Stocks with price impact above the median level are grouped in the “High Price Impact” category (Panel A). Stocks with price impact below the median level are grouped in the “Low Price Impact” category (Panel B).

	(1)	(2)	(3)	(4)
Panel A: High Price Impact				
Log Dollar Value Traded	-0.0839 (-9.80)	-0.0839 (-1.46)	-0.1010 (-27.4)	-0.1010 (-9.04)
Effective Spread (%)	-0.0012 (-1.16)	-0.0012 (-1.35)	-0.0015 (-2.85)	-0.0015 (-3.84)
Realized Spread (%)	0.0046 (6.22)	0.0046 (3.69)	0.0046 (7.34)	0.0046 (4.83)
Price Impact (%)	-0.0062 (-7.37)	-0.0062 (-4.44)	-0.0064 (-10.2)	-0.0064 (-6.96)
Panel B: Low Price Impact				
Log Dollar Value Traded	-0.0379 (-5.29)	-0.0379 (-0.73)	-0.0527 (-20.5)	-0.0527 (-5.67)
Effective Spread (%)	-0.000 (-0.23)	-0.000 (-0.18)	-0.0001 (-0.92)	-0.0001 (-0.38)
Realized Spread (%)	0.0024 (14.9)	0.0024 (3.28)	0.0024 (15.7)	0.0024 (5.24)
Price Impact (%)	-0.0023 (-10.2)	-0.0023 (-3.67)	-0.0024 (-14.1)	-0.0024 (-7.79)
N_{high}	2,990	2,990	2,990	2,990
N_{low}	3,208	3,208	3,208	3,208
T	368	368	368	368
Stock-Venue FE	-	-	X	X
Time FE	-	-	X	X
Error Type	HC1	Clust.	HC1	Clust.

Table 8: Difference-in-Differences Regressions by Exchange

This table reports coefficients (t -statistics) for difference-in-differences regressions for log dollar volume traded, effective spread, realized spread and price impact around fee introduction, split by events. All details are identical to Tables 4, 5 and 6 except that we analyze each event separately. Fee introduction on ARCA is contained in Column (1). Fee introduction on Direct Edge is contained in Column (2). Fee introduction in BATS is contained in Column (3). Regressions in each column contain stock-venue and time fixed effects, and use two-way clustered standard errors (i.e are analogous to Column (4)) in Tables 4 - 6.

	(1)	(2)	(3)
	ARCA	Direct Edge	BATS
Log Dollar Value Traded	-0.0903 (-6.71)	-0.0143 (-0.69)	-0.1171 (-10.3)
Effective Spread (%)	-0.0005 (-0.62)	-0.0011 (-5.57)	-0.0004 (-1.61)
Realized Spread (%)	0.0094 (5.20)	0.0003 (0.55)	0.0007 (2.28)
Price Impact (%)	-0.0112 (-7.75)	-0.0012 (-2.67)	-0.0013 (-3.86)
N	3,026	3,026	3,026
T	116	116	116
Stock-Venue FE	X	X	X
Time FE	X	X	X

Table 9: Difference-in-Differences Regressions — Narrow Window

This table reports coefficients (t -statistics) for difference-in-differences regressions for log dollar volume traded, effective spread, realized spread and price impact around fee introduction, split by events. All details are identical to Tables 4, 5 and 6 except that we use a one month window (approximately 22 trading days) before and after each event.

	(1)	(2)	(3)	(4)
	Model A	Model B	Model C	Model D
Log Dollar Value Traded	0.0106 (0.86)	0.0106 (0.11)	-0.0109 (-2.94)	-0.0109 (-0.76)
Effective Spread (%)	-0.0004 (-0.41)	-0.0004 (-0.44)	-0.0006 (-1.58)	-0.0006 (-3.39)
Realized Spread (%)	0.0013 (2.18)	0.0013 (1.04)	0.0013 (2.62)	0.0013 (2.00)
Price Impact (%)	-0.0014 (-1.94)	-0.0014 (-1.02)	-0.0015 (-3.18)	-0.0015 (-2.61)
N	6,315	6,315	6,315	6,315
T	118	118	118	118
Stock-Venue FE	-	-	X	X
Time FE	-	-	X	X
Error Type	HC1	Clust.	HC1	Clust.