Who Benefits from Securities Exchange Innovation?*

Andriy Shkilko† **Konstantin Sokolov**‡ **Eduard Yelagin**§

May 31, 2022

Abstract. Securities markets continuously innovate to keep pace with technology. It is often debated if such innovation is beneficial, and which market participants capture the benefits. We contribute to this debate by examining liquidity effects of a wide range of proprietary products and services introduced by exchanges in the United States between 2003 and 2017. Exchange innovation is generally associated with liquidity improvements for those investors, who trade in small quantities. The effect is opposite for institutional investors; their trading costs increase, and their market participation declines.

Key words: liquidity, market quality, equity trading, innovation

JEL: G14; G15

*We thank Jim Angel, Tom McInish, Lars Nordén, Mehrdad Samadi, Stefan Voigt, Christian Westheide and conference/seminar participants at the SAFE Market Microstructure Conference, the University of Memphis and the University of St. Andrews for insightful comments. Shkilko acknowledges support of the Social Sciences and Humanities Research Council of Canada and the Canada Research Chairs program. The most current version of this paper may be downloaded from: https://bit.ly/3vqTkxX.

†Wilfrid Laurier University, Canada, e-mail: ashkilko@wlu.ca

‡University of Memphis, United States, e-mail: ksokolov@memphis.edu

§University of Memphis, United States, e-mail: eyelagin@memphis.edu

1. Introduction

The role of securities exchanges in the proper functioning of financial markets is difficult to overstate. Exchanges bring together investors, large and small, allowing them to realize gains from trade, mobilize capital, and incorporate relevant information into prices. To keep pace with the unrelenting march of technology and customer expectations, exchanges continuously innovate. Theoretical models predict that such innovation may have both positive and negative effects on market quality. Likewise, empirical studies find that while some exchange initiatives are beneficial for liquidity, others may be disadvantageous.[1] With such varying findings, it is of interest to ask if exchange innovation is beneficial as a whole, and whether all market participants capture the associated benefits.

We address these questions using a multi-year sample of new product and service offerings by stock exchanges in the United States. These offerings range from technological enhancements in data dissemination to improvements in order processing by exchange engines. The results suggest that exchange innovation is generally associated with lower trading costs for investors who trade in small quantities. In the meantime, institutional investor trading costs increase, and institutional trading volumes decline.

The sample of innovative activities comes from the public record of patents filed by the exchanges. In the U.S., when a company develops new technology it files an application for a patent with the Patent and Trademark Office describing the invention and claiming an exclusive right to it. We collect a sample of all exchange filings in 2003-2017, for a total of 147 patents. The patents capture a wide range of exchange activities. For example, the NYSE patent US-2014089459-A1 describes technology for "high performance data streaming," while Nasdaq patent US-10110540-B2 claims rights to a "message tracking apparatus [... that improves] the latency of a message

[1] See the theoretical models by Menkveld and Zoican (2017), Pagnotta and Philippon (2018), and Cespa and Vives (2022) and emprical studies by Hendershott, Jones, and Menkveld (2011), Hendershott and Moulton (2011), Brogaard, Hagströmer, Nordén, and Riordan (2015), Conrad, Wahal, and Xiang (2015), and Foucault, Kozhan, and Tham (2017)). We discuss these studies and other prior research shortly.

processing system." Like these two examples, the sample patents generally focus on enhancing the speed and efficiency of exchange infrastructure and on improving customer connectivity.

To measure trading costs, we use two data sources, the intraday Trade and Quote (TAQ) database and the Abel Noser institutional trading database. TAQ allows us to compute a set of conventional liquidity metrics such as quoted and effective spreads. The quoted spreads measure displayed liquidity, that is trading costs publicly offered by liquidity providers. In turn, the effective spreads capture liquidity costs that are actually incurred by market participants. Thus, the effective spreads account for the fact that trades are often timed to periods when liquidity is relatively cheap. Both liquidity metrics decline following innovative exchange activities.

What may drive these liquidity cost reductions associated with innovation? To shed light on this question, we examine two components of effective spreads, the price impact and the realized spread. The former component captures adverse selection costs incurred when providing liquidity. The latter reflects liquidity provider inventory and order-processing costs as well as profits. Both components decline post-innovation, consistent with the notion that new exchange technology reduces market maker costs and may also enhance competition for liquidity provision. In addition, the data show that innovation is followed by greater trading volumes, lower price volatility, and greater price efficiency.

Liquidity in the U.S. equity market gradually improves since the early 2000s (Angel, Harris, and Spatt (2015)). As such, it is important to verify that our tests do not merely pick up this background trend. To do so, we de-trend all variables of interest and use only the de-trended variables in regression tests. This approach allows us to attribute the results to exchange innovation with added confidence. In addition, we examine the results in a difference-in-differences (DID) setting against a sample of Canadian securities, for which long-term liquidity trends are similar to those in the U.S. equities. The DID setting further corroborates the notion that the results are driven by exchange innovation. Finally, our findings are robust in the cross-section and replicate for stocks of all sizes, with the strongest effects in large stocks.

TAQ-based liquidity cost proxies discussed so far mainly apply to market participants seeking to trade small quantities of shares. Meanwhile, institutional investors often trade large amounts, and for them the TAQ-based proxies may not be the most suitable (Eaton, Irvine, and Liu (2021)). To shed light on the effect of exchange innovation on institutions, we use Abel Noser data and estimate execution shortfall, a metric commonly used to measure institutional trading costs (e.g., Conrad, Johnson, and Wahal (2001) and Anand, Irvine, Puckett, and Venkataraman (2013)). The results show that exchange innovation is associated with greater execution shortfall. The greater shortfall, in turn, is accompanied by lower institutional trading volume. As such, the data suggest that while exchange innovation may benefit the seekers of small amounts of liquidity, its effects are the opposite for those, who trade large amounts.

What may be behind this dichotomy? We posit that the answer may lie, at least in part, with technologically advanced proprietary trading firms often referred to as high-frequency traders (HFTs). These firms are known for their order submission and information processing speeds, rapid inventory turnover, and generating a lion's share of activity in modern markets. Prior research finds a strong association between high-frequency trading and liquidity provision, yet it also points out that HFTs are highly skilled in avoiding adverse selection and maintaining low inventories (Brogaard, Hagströmer, Nordén, and Riordan (2015), Brogaard, Carrion, Moyaert, Riordan, Shkilko, and Sokolov (2018)). As such, HFT liquidity provision may be beneficial for relatively uninformed investors who trade small quantities, but not necessarily for institutional liquidity seekers.

In addition to liquidity provision, proprietary trading firms regularly engage in liquidity demand. For instance, market maker inventory management as well as many arbitrage strategies are highly time-sensitive and therefore often require that a trader takes liquidity (Brogaard, Hendershott, and Riordan (2014), Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), Boehmer, Li, and Saar (2018), Baron, Brogaard, Hagströmer, and Kirilenko (2019)). The ability of HFTs to quickly harness new technology may give them an advantage over institutions in the race to the outstanding quotes. As such, even though new exchange technology becomes available to everyone simultaneously, not all market participants may be sufficiently equipped to use it. As a result, their liquidity costs may increase, and their market participation may decrease.

Our findings contribute to an ongoing discussion of modern market structure and particularly the securities exchange industry. When the U.S. exchanges demutualized and became public companies in the early 2000s, their decision making has expectedly shifted towards the profit motive. Market observers often allege that this motive may not align with the interests of all exchange customers and of the society at large. For instance, former SEC Commissioner Robert Jackson suggests that "exchanges [...] have developed puzzling practices that look nothing like the competitive marketplaces investors deserve."[2] Similarly, Spatt (2020) argues that much of U.S. market structure is characterized by exchange market power, potentially leading to anti-competitive practices.

These concerns reflect in recent industry trends, whereby institutional investors often support launching new trading venues. The IEX, also known as the Investors Exchange, began operating as an official exchange in 2016 aiming to keep HFT participation to a minimum and instead focussing on institutional investors. Another addition to the trading landscape, the Members Exchange (MEMX) came online in 2020 and sets itself apart by focusing on the needs of its institutional members rather than shareholder profits.[3]

Our results suggest that the premises, on which these markets are founded may have merit. Even though exchange innovation by the incumbents likely does not have nefarious intentions, it may inadvertently benefit market participants, who are best equipped for technological change. In this regard, establishing markets that cater to the needs of the institutional community may be socially optimal. This logic echoes Biais, Foucault, and Moinas (2015), who model interactions

[2] R. Jackson Jr., "Unfair Exchange: The State of America's Stock Markets," September 19, 2018 (`https://bit.ly/3e5LHpK`).

[3] K. Dew, "Is There an Answer to MEMX?" The Tabb Forum, July 4, 2020 (`https://bit.ly/3gtMAeA`) and P. Stafford and N. Bullock, "Wall Street heavyweights back new exchange rival to NYSE, Nasdaq," Financial Times, January 7, 2019 (`https://on.ft.com/2EBPxt2`).

between investors that operate at different trading speeds and propose segmenting markets into those that cater to fast and slow traders.

2. Related literature

Several theoretical models consider the effects of advancements in exchange technology and predict that such technology may be both beneficial and detrimental. Budish, Cramton, and Shim (2015) argue that continuous order processing used by the majority of modern exchanges may exacerbate adverse selection costs and therefore harm liquidity. Menkveld and Zoican (2017) show that increasing exchange engine processing speed may have both positive and negative effects on liquidity. In their model, the direction of the effect may vary over time and in the cross-section and depends on the mix of liquidity traders and news events. Pagnotta and Philippon (2018) show that in a competitive environment increasing exchange speed may lead to lower trading fees and greater investor participation. They however caution that purely technological improvements to trading technology may lead to limited welfare gains.

Perhaps the closest to our study, Cespa and Vives (2022) model an environment, in which exchanges supply technology services to liquidity providers. The fast adopters of new technology use it to improve liquidity, allowing other market participants to hedge more aggressively for a positive welfare effect. As a side effect, liquidity providers who are slow to adopt new technology face high competitive pressures and lower payoffs. Our results echo the implications of this model. While we find that exchange innovation tends to benefit liquidity consumers, we also show that it results in greater trading costs and lower market participation by institutional investors, who are generally slower to adopt new technology.

The empirical literature examines a number of technological enhancements to the exchange trading process. The outcomes of such enhancements vary. For instance, Hendershott, Jones, and Menkveld (2011) show that exchange-driven automation of liquidity provision results in liquidity improvements. Conversely, Hendershott and Moulton (2011), Foucault, Kozhan, and Tham (2017), and Indriawan, Pascual, and Shkilko (2021) find that exchange technology enhancements that benefit liquidity-demanding strategies, including various arbitrage strategies, have adverse liquidity effects. Brogaard, Hagströmer, Nordén, and Riordan (2015) and Conrad, Wahal, and Xiang (2015) study co-location, a practice whereby trading firms are allowed to place trading servers near exchange matching engines. They report that the practice is mainly used by liquidity suppliers and therefore benefits liquidity. As such, empirical findings appear to vary across innovative technologies and practices, making it difficult to make statements about exchange innovation as a whole. We therefore posit that our analysis of an extended sample of innovative exchange offerings may bring the literature one step closer to understanding the general effects of exchange innovation.

Our study is related to the vast literature that examines liquidity costs and to a sub-set of this literature that focuses on institutional investors. Anand, Irvine, Puckett, and Venkataraman (2013) suggest that despite the overall trend to lower liquidity costs in the 21st century, institutional investor trading costs did not decline. This result echoes earlier findings by Goldstein and Kavajecz (2000) and Jones and Lipson (2001), who show that regulatory actions that reduce spreads may not have similar liquidity-boosting effects for institutional traders. Eaton, Irvine, and Liu (2021) generalize this finding to caution that changes in institutional trading costs do not always align with changes in conventional TAQ-based liquidity metrics.

Our arguments are also rooted in an extensive literature on automated and high-frequency trading. A number of studies in this literature show that HFTs perform several important functions, among which liquidity provision and arbitrage stand out as dominant (Brogaard, Hendershott, and Riordan (2014), Chaboud, Chiquoine, Hjalmarsson, and Vega (2014), Brogaard, Hagströmer, Nordén, and Riordan (2015), Boehmer, Li, and Saar (2018)). In the meantime, HFTs are highly skilled at avoiding adverse selection and maintaining low inventories. Therefore, their liquidity supply is likely more beneficial to investors seeking small amounts of liquidity than to institutional investors (Brogaard, Hagströmer, Nordén, and Riordan (2015), Brogaard, Carrion, Moyaert, Riordan, Shkilko, and Sokolov (2018)). Two recent studies show explicitly that HFTs do not always benefit institutional investors (Korajczyk and Murphy (2019), van Kervel and Menkveld (2019)).

Prior research suggests that changes in financial markets may have different effects in the cross-section of stocks. For instance, Hendershott, Jones, and Menkveld (2011) show that automation of market making has less of an effect on small stocks than on large stocks. They ascribe this finding to the possibility that algorithms are less commonly used in small stocks. Haslag and Ringgenberg (2020) find that market fragmentation is associated with liquidity improvements in large stocks and worse liquidity in small stocks. Foley, Liu, Malinova, Park, and Shkilko (2020) argue that modern liquidity providers tend to focus on large and frequently traded securities and suggest that liquidity in smaller securities may be improved by introducing additional market making incentives. Given these findings, we too ask if the effects of exchange innovation differ in the cross-section. The data show that the effects are consistent for stocks of all sizes, and are occasionally stronger for the large stocks.

3. Data and metrics

Our data come from three main sources: (i) a public depository of patents maintained by Google, LLC (`https://patents.google.com`), (ii) the Trade and Quote (TAQ) database provided by the New York Stock Exchange, and (iii) the Abel Noser (formerly, ANcerno) institutional trading database. For the difference in differences (DID) tests, we also use a fourth data source – Canadian data on liquidity costs and trading volume. These data are from the Canadian Financial Markets Research Centre (CFMRC) (`https://bit.ly/3czbKUV`).

3.1 Patent data

From Google, we obtain 147 patents that are filed by the U.S. exchanges (e.g., the NYSE, Nasdaq, BATS) between 2003 and 2017. These patents cover a broad range of exchange activities, with a general focus on optimizing data dissemination and order management. Even though the filers use rather arcane language to describe their inventions, in most cases it is possible to infer the purpose of the proposed technologies. For example, and as we mentioned earlier, the NYSE patent US-2014089459-A1 describes technology for "high performance data streaming", while the Nasdaq patent US-10110540-B2 claims rights to a "message tracking apparatus [... that improves] the latency of a message processing system." Like these two examples, the majority of patents describe incremental enhancements to information transmission and order processing speeds.

We note that while the patents in our sample capture a wide range of exchange activities, they do not represent an exhaustive list of new products and services. For instance, during our sample period exchanges start offering co-location – a practice that allows trading firms to place their servers as close as possible to the exchange engine. Co-location is a concept rather than a technology that can be patented, so our sample of events does not include the introduction of co-location. Rather, the sample contains products or services that, among other things, may improve data dissemination and trader connectivity within the co-location centers. As such, instead of capturing momentous yet rare market structure changes, the event sample contains numerous incremental adjustments to exchange technology.

Figure 1 reports monthly time series of exchange patent filings. In most months, the number of patents varies between zero and one, but there are also months with greater numbers of filings, with a maximum of nine patents filed in May 2005. In subsequent tests, we account for the number of patents filed in close succession and report the results on a per-patent basis.

[Figure 1]

In the United States, the prevailing corporate practice is to file a provisional patent application prior to deploying a new product, service, or technology. Filing after deployment jeopardizes intellectual property rights since the World Intellectual Property Organization considers inventions unpatentable if they have fallen into the public domain prior to filing.[4] Since product and service implementation involves informing exchange customers and regulators, and therefore brings inventions into the public domain, exchanges file provisional applications as the first step of the implementation process. We use the dates of such filings, also known as *patent priority dates*, to proxy for event dates in subsequent analyses.

Even though it is common to deploy new products soon after filing a provisional application, we do not know the exact dates of deployment of new products and services. To stay conservative, we define the three-month window surrounding the priority date as *the implementation period*. As the illustration below shows, in the main analysis we compare market quality in the pre-implementation period (the three months prior to the implementation period) and the post-implementation period (the three months following the implementation period). In the robustness section, we examine several alternative event window lengths, including shorter implementation periods, and find similar results.

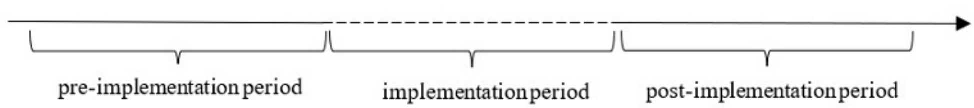

3.2 Liquidity metrics

We obtain liquidity metrics from two sources; the NYSE TAQ database and the Abel Noser database. Using TAQ, we compute quoted and effective spreads as well as the two components of the latter, that is, price impacts and realized spreads. The *quoted spread*, or the National Best Bid and Offer (NBBO), measures displayed liquidity and is computed as the difference between

[4] See Rule 64 of the Patent Cooperation Treaty, to which the U.S. is a party: https://bit.ly/3eMqRM6.

the national best offer and the national best bid. Further, to measure trading costs incurred by the liquidity demanders, we compute the *effective spread* as twice the signed difference between the traded price and the quote midpoint at the time of the trade. Next, to assess the levels of adverse selection, we compute the *price impact* as twice the signed difference between the quote midpoint at the time of the trade and the midpoint at a future time. Finally, the *realized spread* is computed as the difference between the effective spread and price impact and is often associated with liquidity provider inventory and order processing costs as well as profits (e.g., Hendershott, Jones, and Menkveld (2011), Brogaard, Hagströmer, Nordén, and Riordan (2015)).

To sign trades, we rely on the Lee and Ready (1991) algorithm. Chakrabarty, Pascual, and Shkilko (2015) show that this algorithm performs well in modern markets. All variables are scaled by the corresponding quote midpoints. When computing daily aggregates, we weight quoted spreads by the time they are outstanding, and we weight all trade-related metrics by the corresponding trading volume. To compute price impacts, we use five-minute horizons because our sample goes back to the period when trading was slower than it is today. During this period, using such horizons is conventional (e.g., Hendershott, Jones, and Menkveld (2011)), and we use them through the entire sample period for consistency. In a recent study, Shkilko and Sokolov (2020) show that horizon length does not qualitatively affect inferences based on price impacts.

The above-mentioned TAQ-derived metrics are generally considered representative of small investor trading costs. This is because the best bid and ask quotes are usually valid for several hundred or several thousand shares, representing quantities sufficient to execute small trades, but not the trades of large institutional investors. To estimate liquidity metrics for such larger investors, we use the Abel Noser database. Puckett and Yan (2011) show that even though this database does not contain trading by all institutions, the activity it captures is representative of that by an average institution.

To proxy for institutional trading costs, we use the *execution shortfall* metric of Conrad, Johnson, and Wahal (2001) and Anand, Irvine, Puckett, and Venkataraman (2013). The metric is computed as the signed difference between the volume-weighted execution price of an institutional order in stock i on day t and the opening price for the stock that day. Abel Noser differentiates between institutional buys and sells, so no additional trade signing is required. As with all trading cost metrics, we volume-weight execution shortfall.

3.3 Price efficiency metrics

In addition to examining the effects of exchange innovation on liquidity costs, we also measure its effects on price efficiency. To do so, we use three standard efficiency metrics: the *variance ratio* of Lo and MacKinlay (1988), *return autocorrelation* as in Hendershott and Jones (2005), and the *price delay* of Hou and Moskowitz (2005).

The first metric, the variance ratio, relies on the notion that if prices follow a random walk, return variance should increase linearly in return horizon. With this in mind, for each stock-day we compute the absolute difference between one and the ratio of the variance of $k \times l$-second midpoint returns to k times the variance of l-second midpoint returns:

$$|1-VR| = \left|1 - \frac{\sigma^2_{k\times l}}{k \times \sigma^2_l}\right|, \tag{1}$$

where $(l,\ k\times l)$ is $(15,60)$ or $(60,300)$. The closer the variance ratio is to zero, the more prices resemble a random walk implying greater efficiency. We choose $k \times l$ of 300 seconds to be consistent with the horizon of price impact estimates. As we mention before, during the early years of our sample period the market is slower, and longer estimation horizons are appropriate. To account for greater speeds that prevail later in the sample period, we follow Comerton-Forde, Grégoire, and Zhong (2019) and also examine $k \times l$ of 60 seconds. In the Appendix, we report results for alternative estimation horizons.

The second metric, return autocorrelation, makes use of the notion that in a frictionless market prices should be unpredictable, and as such midpoint returns should have zero autocorrelation.

It is computed for each stock-day as:

$$autocorrelation = corr\left(ret_{l,\ t}, ret_{l,\ t-1}\right), \tag{2}$$

where $ret_{l,\ t}$ is the t^{th} midpoint return at horizon l. For consistency with the metrics discussed earlier, in the main results we focus on return autocorrelations computed at the 15-second and 300-second horizons, with additional horizons reported in the Appendix. Smaller autocorrelation estimates suggest greater efficiency.

The third metric, the price delay, assumes that efficient prices should instantly incorporate public market information. Accordingly, lagged market returns should have no predictive power for individual stocks returns. To compute this metric, we begin by running the following regression for each stock-day:

$$ret_l = \alpha + \beta ret_{m,\ l} + \sum_{\tau=1}^{10} \gamma_\tau ret_{m,\ l-\tau} + \varepsilon_l, \tag{3}$$

where ret_l is the quote midpoint return during time interval l, and $ret_{m,\ l}$ is the return on S&P 500 proxied by the SPY ETF. For consistency, we use the same frequencies for l as we did when computing the other price efficiency metrics, with additional horizons reported in the Appendix. We then define the R^2 from regression (3) as unconstrained, R^2_u. Next, we estimate regression (3) without the lagged market returns, effectively constraining γ to zero, and define the corresponding R^2 as constrained, R^2_c. Finally, for each stock-day, we compute:

$$price\ delay = 1 - \frac{R^2_c}{R^2_u}. \tag{4}$$

A smaller delay suggests greater efficiency.

3.4 Sample

To define the sample, we begin with all common equities listed on the NYSE, Nasdaq, or AMEX at the beginning of the sample period, January 2003. We then drop the stocks that do not survive the entire 2003-2017 period and from the remaining equities select 900 stocks with the greatest market capitalization. Our access to the Abel Noser dataset ends in 2013, so the sample period for institutional analyses is shorter than that for the TAQ analyses. In the robustness section, we show that the TAQ results do not change if we use the shorter 2003-2013 period that matches the Abel Noser period.

Table 1 reports the summary statistics. The average stock has about $12 billion in market capitalization and trades close to 2.5 million shares a day at a price of $47. From TAQ, its quoted and effective spreads are 15.84 and 14.38 bps, with price impacts of 7.35 bps and realized spreads of 5.57 bps. In turn, daily institutional volume in Abel Noser data is 0.22 million shares per day (about 9% of total volume), and the typical implementation shortfall is 13.15 bps. Both these figures are consistent with earlier studies that use Abel Noser data.

[Table 1]

We note the cross-sectional variation in many of the above-mentioned variables that should be expected from a sample of 900 equities. For instance, while the average sample firm has market capitalization of $11.74 billion, the median firm is smaller, with a capitalization of $2.92 billion. Trading costs also vary in the cross-section. For example, the effective spreads are 6.7 bps in the 25th percentile of firms and 18.06 bps in the 75th percentile. We observe similar variation for price efficiency metrics, with 15-second return autocorrelation for instance ranging from 0.056 to 0.064. In a later section, we show that the main results are generally preserved in the cross-section when we separately examine large, medium, and small firms.

3.5 Adjusting for the long-term trend and other possible confounders

The event windows in this study span nine months around patent priority dates. As such, our analyses capture medium-term changes. We note that at the background of these changes are the long-term trends in liquidity and volume. Angel, Harris, and Spatt (2015) report that in the 21st century spreads generally trend downward, while trading volume trends upward. To mitigate the effects of these trends, we regress every variable of interest on a time trend as follows:

$$DepVar_{it} = \alpha + \beta Trend_t + \varepsilon_{it}, \tag{5}$$

and use the estimated residuals $\hat{\varepsilon}_{it}$ as dependent variables in all subsequent tests.

For illustration, Figure 2 plots four of the resulting detrended series: effective spread, execution shortfall, total trading volume, and institutional volume. The plots show that the above-mentioned long-term trends in the variables of interest are successfully removed by the procedure in equation (5). As such, we are reasonably confident that the changes identified in the subsequent tests result from exchange innovation rather than the background trends.

[Figure 2]

To further assuage concerns about possible confounding background processes, in a subsequent section we also introduce a DID analysis that compares the variables of interest from U.S. data to their Canadian equivalents. Our Canadian data are daily and resemble those available through CRSP. As such, the DID analyses rely on spread proxies proposed by Corwin and Schultz (2012) and Abdi and Ranaldo (2017) as well as on volume figures.

4. Empirical results

4.1 TAQ-based trading costs

We begin the analysis by examining the effects of exchange innovation on liquidity metrics derived from TAQ. Figure 3 provides an illustration of these effects focusing on quoted and effective spreads in a simple event study setting. We report daily averages of the spread metrics in the pre-innovation period and the post-innovation period as well as the average for each period. In this test, the variables are not yet de-trended. Quoted spreads decline from 15.54 bps in the pre-innovation period to 14.67 bps in the post-innovation period, a 0.87 bps change. Meanwhile, effective spreads decline by 0.92 bps. We note that the average innovation period contains 4.55 patents, so the above-mentioned reductions are 0.19 and 0.20 bps per patent. It is important to point out that these results do not yet account for the long-term trend in liquidity costs and for the determinants of these costs, that is volume and volatility. We discuss these controls and the associated results in the regression setting that follows.

[Figure 3]

Hendershott, Jones, and Menkveld (2011) and Brogaard, Hagströmer, Nordén, and Riordan (2015) show that volume and volatility are first-order trading cost determinants that should be controlled for in event studies like ours. We account for the effects of these variables in the following regression framework:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it}, \tag{6}$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing quoted or effective spread for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. The model controls for stock fixed effects. If more than one patent is filed on day t, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent.

Table 2 reports univariate (Panel A) and regression results (Panel B) quantifying innovation-related changes in quoted and effective spreads. Both variables decline post-implementation. In the regression setting without controls (specifications 1 and 3), implementation leads to lower quoted and effective spreads with statistically significant declines of, respectively, 0.011 and 0.009 standard deviations. These figures translate to spread reductions of 0.068 and 0.064 bps per patent.[5] We note that the magnitude of these changes is somewhat smaller than that of the univariate figures, which is likely attributable to the fact that we de-trend the variables of interest.

[Table 2]

Specifications 2 and 4 include volume and volatility controls. The economic significance of the *Post* coefficients declines further, to 0.008 and 0.006 standard deviations, representing quoted and effective spread reductions of 0.050 and 0.048 bps. At first glance, the economic significance of these estimates may appear trivial, but we note that it represents the change per patent. Taken together, the cumulative effect of all patents implemented during the sample period amounts to a 7.35 bps decline in quoted spreads and a 7.06 bps decline in effective spreads or about one half of the sample averages for these statistics – a rather substantial change.

When interpreting the cumulative effects, it is important to remember that exchange innovation does not occur in a vacuum. While innovation may reduce trading costs, other changes to market structure such as the automation of market order submissions (Hendershott and Moulton (2011)), the expansion of direct market access (Chakrabarty, Jain, Shkilko, and Sokolov (2021)),

[5] To obtain these figures, we use standard deviations computed across days rather than across stocks as in Table 1. Such standard deviations better align with the coefficient estimates from the event-study regressions that constitute the bulk of our analyses. The standard deviations used for economic interpretation are reported in Table A1 of the Appendix.

the emergence of dark trading (Zhu (2014), Comerton-Forde and Putniņš (2015)), and the proliferation of latency arbitrage (Shkilko and Sokolov (2020), Aquilina, Budish, and O'Neill (2021)) have been shown to increase transaction costs. These effects are likely to partly offset those of innovation. As such, exchange innovation should be viewed as one of many determinants of liquidity costs that often offset each other in a complex system of interactions.

4.2 Trading cost components

In competitive markets, a decline in effective spreads is usually caused by a reduction in the cost of liquidity provision. Researchers typically examine two components of this cost, those related to adverse selection and inventory holdings. We too ask how these components are affected by exchange innovation by decomposing the effective spread into price impact and realized spread. The former reflects the cost incurred by liquidity providers due to trading against informed order flow (adverse selection), and the latter captures (i) the remaining costs, including those associated with order processing and inventory holdings, as well as (ii) liquidity provider profits.

Table 3 reports that both price impacts and realized spreads decline post-implementation, respectively, by 0.11 bps and 0.46 bps in the univariate setting and by 0.001 and 0.003 standard deviations in the regression setting with volatility and volume controls. As we mention previously, these figures are per patent. The former result is consistent with the notion that exchange innovation generally allows liquidity providers to reduce exposure to being picked off by informed traders. The interpretation of the latter result is somewhat more nuanced. On the one hand, it may suggest that innovation helps liquidity providers to better manage order-processing and inventory costs. On the other hand, it may indicate that a typical round of innovation leads to greater competition among liquidity providers reducing their profits. Our data do not allow us to examine which of these effects dominates, yet we believe that both of them may be at play.

[Table 3]

Taken together, the event study results suggest that exchange innovation is associated with lower trading costs, which are potentially attributable to reductions in liquidity provider costs and profits. We emphasize that these results mainly reflect trading cost of relatively small investors. We examine institutional trading costs shortly.

4.3 Difference in differences

One downside of the events that we study is that they affect all sample stocks at once. It is therefore possible, although unlikely, that another series of events occurring simultaneously with exchange innovation may confound the results. When it comes to this concern, one mitigating factor is that the sample includes multiple patent filings that are spread across multiple years. As such, and aside from the long-term trend accounted for via de-trending in equation (5), it is difficult to think of recurring events that are sufficiently persistent to confound the findings.

This said, to further assuage concerns about possible confounding effects we supplement the earlier results with a DID setup. To do so, we use a control sample of stocks that trade in Canada matched to the main U.S. sample by market capitalization and trading volume. Firm sizes and trading volumes routinely change during our multi-year sample period. To keep up with these changes, we re-match each sample stock with the most suitable control at the beginning of each calendar year. For instance, in 2012 we match the stocks by their average market capitalization and total trading volume in 2011.

As we mention previously, our Canadian data have daily granularity and resemble CRSP data. As such, to compare the two samples we rely on three low-frequency spread proxies developed for CRSP-like datasets: the end-of-day quoted spread, EOD, the Corwin and Schultz (2012) effective spread proxy, CS, and the effective spread proxy developed by Abdi and Ranaldo (2017), AR. Abdi and Ranaldo (2017) show that the EOD quoted spread is the most accurate low-frequency liquidity proxy. Since our high-frequency metrics distinguish between quoted and effective spreads, we opt to complement the *EOD* quoted spreads with the two above-mentioned effective spread metrics for consistency.

For the DID analysis, we use the differences between the sample stocks and their controls in a regression setup similar to that discussed earlier:

$$\Delta DepVar_{ijt} = \alpha_{ij} + \beta_1 Post_t + \beta_2 \Delta Volume_{ijt} + \beta_3 \Delta Volatility_{ijt} + \varepsilon_{ijt}, \quad (7)$$

where $\Delta DepVar_{ijt}$ is the difference between the spread proxies (i.e., *EOD*, *CS*, and *AR*) computed on day t for a pair of stocks that includes a U.S. sample stock i and its matched Canadian counterpart j, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, and $\Delta Volume_{ijt}$ and $\Delta Volatility_{ijt}$ controls are the differences between the volume and volatility estimates for each pair of stocks i and j on each day t. If more than one patent is filed during the implementation period, *Post* equals to the number of patents. This approach allows us to identify the effect per patent.

Table 4 confirms the earlier results in that exchange innovation is followed by a decline in both quoted and effective spreads. The coefficient estimates of the *Post* dummy variable are statistically significant with a negative sign across all specifications. These results allow us to suggest with added confidence that liquidity improvements documented thus far are driven by innovative activities of U.S. exchanges.

[Table 4]

4.4 Volume and volatility

In the earlier sections, we use volume and volatility as regression controls since these two variables are known trading cost determinants. Now, we ask if these two variables themselves undergo changes as a result of exchange innovation. The exchanges derive a substantial portion of their revenue from trading volume, and as such we expect their innovative activity to facilitate increases in volume. In addition, lower trading costs identified in the previous section as a consequence of innovative activity may themselves result in greater trading volume as more agents engage in trading activity when doing so is cheaper.

When it comes to the second variable of interest – volatility, we rely on theory and prior empirical evidence to form expectations. Specifically, Roşu (2019) suggests that trading strategies that generate adverse selection often result in greater volatility. Shkilko and Sokolov (2020) and Indriawan, Pascual, and Shkilko (2021) empirically confirm this association. Since our data show that adverse selection declines post-innovation, we expect that volatility may also decline.

The results in Table 5 examine these issues using simple differences (Panel A) and the equation (6) framework (Panel B). As previously, the variables are detrended. In the regression setting, the data show that volume increases by 0.006 standard deviations per patent, while volatility decreases by 0.010 standard deviations per patent. These results suggest that innovation is associated with greater price stability and also potentially with greater gains from trade, as more agents may come to the market to participate in asset exchange. The DID analysis comparing U.S. volume and volatility statistics to their Canadian counterparts corroborate these results (Panel C).

[Table 5]

4.5 Exchange innovation and institutional traders

The reductions in spreads documented in the previous section suggest that stock exchange innovation may improve liquidity for small trades. Modern markets however serve a diverse range of market participants, and smaller spreads may not necessarily imply lower execution costs for all of them, particularly the institutional traders who often trade large amounts. Along these lines, Eaton, Irvine, and Liu (2021) suggest that even though the bid-ask spreads may decrease following certain market reforms, institutions may not always capture the benefits of such decreases. To re-examine this issue in the exchange innovation setting, we take a closer look at institutional trading volume and execution shortfall.

Table 6 shows that institutions incur greater transaction costs after exchanges innovate. On average, the post-innovation period is associated with a 0.001 standard deviations, or 0.010 bps, increase in execution shortfall per patent. Notably, the data also suggest that the increase in execution shortfall is accompanied by a decline in institutional trading volume, by 0.001 standard deviations per patent. As such, the increase in transaction costs that follows exchange innovation appears to result in institutions trading less than before.

[Table 6]

Even though new exchange technology becomes available to everyone at the same time, its benefits are likely to accrue mainly to those market participants, who have advanced technological capabilities. In modern markets, such participants are often referred to as high-frequency traders (HFTs), and liquidity provision is one of their main activities (Boehmer, Li, and Saar (2018)). The market structure literature suggests that HFTs are able to avoid adverse selection and manage inventories more effectively than non-HFTs, and as such liquidity provided by them is often cheaper (Hoffmann (2014), Brogaard, Hagströmer, Nordén, and Riordan (2015), Jovanovic and Menkveld (2016), Chordia, Green, and Kottimukkalur (2018)).

Our results are consistent with the notion that new exchange technology tends to benefit liquidity providers, as we find an association between innovation and lower spreads. The data also suggest that the channels through which exchange innovation facilitates liquidity provision is the reduction in adverse selection and possibly other costs of market making such as inventory costs. Still, one of the features of HFT inventory management is their preference for small inventory accumulations, which is often associated with quoting relatively small amounts of liquidity at any given time (Brogaard, Carrion, Moyaert, Riordan, Shkilko, and Sokolov (2018)). Together with adverse selection avoidance, this preference implies that HFT liquidity provision may be more beneficial to relatively uninformed investors who trade small quantities than to the institutional liquidity seekers. The data corroborate this possibility and show that exchange innovation is associated with greater institutional trading costs.

4.6 Price efficiency

Modern markets achieve price efficiency in two ways. First, liquidity providers continuously monitor their quotes and adjust them to reflect the most up to date information on asset supply and demand as well as information on prices of correlated assets. Second, arbitrageurs aim to quickly pick off market maker quotes that have become stale, thereby forcing prices to adjust. The marketable orders used in the process of such arbitrage typically generate adverse selection (Foucault, Kozhan, and Tham (2017)).

Our earlier results show that adverse selection declines post-innovation, suggesting that new exchange technology allows market makers to monitor and update their quotes more effectively. On the one hand, such improved quote management may simply change the balance of responsibilities for price adjustments between arbitrageurs and market makers with no effect on the overall level of price efficiency. On the other hand, new exchange technology may facilitate the very process of price adjustment and result in improved price efficiency, especially at shorter time intervals where such technology is likely the most beneficial.

We examine these possibilities by estimating equation (6) for the three price efficiency metrics discussed earlier: the variance ratio, return autocorrelation, and the price delay. The results in Table 7 are consistent with the notion that exchange innovation is often associated with price efficiency improvements. For instance, the variance ratios computed for the (15 sec, 60 sec) interval decline by 0.13% per patent, whereas the 15-second autocorrelations decline by 0.15% per patent. We note that not all metrics show statistically significant changes. The variance ratios and return autocorrelations computed over longer intervals do not significantly change post-innovation. These results are consistent with our earlier suggestion that new technology may be the most useful for price efficiency at shorter time horizons. In the Appendix, we report the results from alternative estimation horizons and confirm that most price efficiency improvements occur at the shorter end of the estimation horizon spectrum.

[Table 7]

4.7 The cross-sectional effects

Prior research suggests that changes in financial markets may have different effects in the cross-section of stocks. For instance, Hendershott, Jones, and Menkveld (2011) show that automation of market making has less of an effect on liquidity in small stocks compared to large stocks. They ascribe this finding to the possibility that algorithms may be less commonly used in small stocks. Haslag and Ringgenberg (2020) find that market fragmentation is associated with liquidity improvements in large stocks and worsening of liquidity in small stocks. Foley, Liu, Malinova, Park, and Shkilko (2020) argue that modern liquidity providers tend to focus on large and frequently traded securities and suggest that liquidity in less frequently traded securities may be improved by introducing additional market making incentives. Given these findings, it may be of interest to examine whether the effects of exchange innovation too differ in the cross-section.

To do so, we split the sample into terciles by market capitalization, with each tercile containing 300 stocks. We then estimate the following regression model:

$$\begin{aligned} DepVar_{it} = {} & \alpha_i + \gamma_t + \beta_1 Post_t + \beta_2 Post_t \times Medium_{it} + \beta_3 Post_t \times Large_{it} + \\ & \beta_4 Volume_{it} + \beta_5 Volatility_{it} + \varepsilon_{it}, \end{aligned} \tag{8}$$

where *Medium* and *Large* are dummy variables for the medium and large stocks, and all other variables are previously defined. In this setup, the *Post* dummy absorbs the effect for the small stocks, and the two interaction variables capture the incremental effects for the other two size categories, that is the effects *in addition to* the small stock effect.

Table 8 shows that virtually all findings reported earlier for the full sample are observed in the cross-section, with some results amplified in larger stocks. For instance, specification [1] shows that quoted spreads decline in stocks of all sizes, but the decline is the greatest in large stocks, with the total effect captured by the combination of *Post* and *Post* $\times$ *Large* coefficients being $-0.010(=-0.005+(-0.005))$. In the meantime, trading costs proxied by effective spreads decline similarly across all stocks as indicated by the significant *Post* coefficient and insignificant *Post* $\times$ $Medium$ and *Post* $\times$ *Large* coefficients (specification [2]).

[Table 8]

It is interesting to think why quoted spreads decline more in large stocks than in small stocks, yet effective spreads decline similarly across all stock sizes. Liquidity timing stands out as one possible reason. The relative abundance of liquidity in large stocks may allow traders to better time liquidity demand even when the spreads are not at their narrowest. As such, improvements in quoted spreads may have a smaller effect on realized trading costs in large stocks. Meanwhile, in small stocks liquidity timing is more challenging, and improvements in displayed liquidity may therefore reflect more readily in lower trading costs. Going further through the set of market quality statistics in Table 8, similar patterns arise. For instance, specifications [5] and [6] show that volume increases equally for stocks of all sizes, while volatility declines in all stocks, with a slightly greater decline in large stocks.

One result that deserves additional discussion is the interplay between execution shortfall and institutional trading volume in specifications [7] and [8]. On the one hand, the data suggest that execution shortfall increases post-innovation in medium and large stocks, but not in small stocks. On the other hand, institutional volume appears to decline equally across all stock size categories.

While the results for medium and large stocks are consistent with the earlier discussion, it may not be immediately clear why institutional trading volume in small stocks declines without a corresponding increase in execution shortfall. We suggest two possible reasons for this result. First, statistical power issues may be at play when we measure execution shortfall in small stocks. As such, this particular result should be interpreted with caution. Second, it is possible that the result is driven by spillover effects. Namely, greater trading costs in large and medium stocks may affect institutional portfolio trading decisions across the entire stock size spectrum, thus affecting stocks of small sizes in addition to those of larger sizes.

Finally, when it comes to price efficiency the results in Table 9 are again consistent with the earlier findings in that exchange innovation generally improves price efficiency, and the improvements tend to concentrate at shorter estimation horizons. The effects at such shorter horizons are also more pronounced for the large stocks. For instance, while the *Post* coefficient for the variance ratio computed for the (15s, 60s) interval is -0.001 for the small stocks, suggesting post-innovation improvement, the composite coefficient for the large stocks is -0.003 (= -0.001-0.002) pointing to an even greater improvement.

[Table 9]

4.8 Robustness

In an earlier section, we mentioned that while TAQ data are available to us between 2003 and 2017, the Abel Noser data end in 2013. Our results are however robust to shortening the TAQ sample period to match the Abel Noser sample period. We also mentioned that our results are robust to varying the length of event windows, which in the main sample include a three-month pre-implementation period, a three-month implementation period, and a three-month post-implementation period. In Table 10, we report the results of these robustness checks, with line 1 referencing the base regression case, that is years 2003-2017 and the 3-3-3 event window.

[Table 10]

First, we ask if the main results hold if we shorten the TAQ sample period to 2003-2013 to match it with the Abel Noser sample period. The results reported in line 2 of Table 10 confirm those reported earlier. More specifically, the data show that quoted, effective, and realized spreads all decline subsequent to innovative activities by exchanges. Price impacts and price volatility also decline, while trading volume increases. Next, in line 3 we ask if the results change if we omit the financial crisis years, that is 2008 and 2009. These two years are characterized by substantial volatility and illiquidity, and as such we would like to ensure that the results are robust to excluding them. All estimated coefficients remain statistically significant, and their economic significance does not change notably compared to the base case.

In the last set of checks, we revert to the main sample period of 2003-2017 and examine several adjustments to event windows. As we mention earlier, our data do not contain exact innovation implementation dates, although it is common practice to implement a patented innovation soon after filing a provisional patent application. In the main sample, we allow for a relatively wide implementation window of three months to accommodate possible implementation delays. In lines 4 and 5 of Table 10, we reduce the length of the implementation period first to two and then to one month. In turn, in lines 6 and 7 we expand the length of the pre- and post-implementation periods to four and five months. All of the results remain statistically significant and are in line with the base case.

5. Conclusion

Operating in a competitive and technologically advanced environment, securities exchanges continuously innovate, offering new products and services to its customers. Theoretical models predict that such innovation may have both positive and negative effects on market quality. Examining individual cases of innovative exchange activities, empirical studies find that some exchange initiatives are beneficial for liquidity, whereas others may be disadvantageous. In this study, we ask if exchange innovation is beneficial as a whole, and if all market participants capture the associated benefits.

To answer this question, we use a multi-year sample of innovative product and service offerings by the U.S. stock exchanges. These offerings range from technological enhancements in data dissemination to improvements in order processing by exchange engines. The results suggest that exchange innovation is generally associated with lower trading costs for investors who trade in small quantities. In the meantime, institutional investor trading costs increase, and market participation by institutions declines.

We ascribe this result to the unique ecosystem that prevails in most modern markets. Proprietary firms commonly known as HFTs play a dominant role in supplying liquidity, but also consume substantial amounts of liquidity in the process of inventory management and arbitrage. These firms are more technologically nimble than other market participants and therefore are able to capture the benefits of exchange innovation more successfully than others. Since liquidity provided by such firms is known to benefit investors seeking to trade small rather than large quantities, the effects of exchange innovation vary in trader size.

We note that our results do not necessarily imply that the nature of exchange innovation is nefarious. Neither do we suggest that exchanges cater to certain groups of market participants over others. Instead, the findings are consistent with the notion that in a highly technologically advanced industry, some participants will lag behind in their ability to benefit from technological progress. In this light, theoretical models that suggest that investors may benefit from separating venues into fast and slow may have merit.

References

Abdi, F., and A. Ranaldo, 2017, "A simple estimation of bid-ask spreads from daily close, high, and low prices," *Review of Financial Studies*, 30(12), 4437–4480.

Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2013, "Market crashes and institutional trading: Evidence from US equities during the financial crisis of 2007-08," *Journal of Financial Economics*, 108, 773–797.

Angel, J. J., L. E. Harris, and C. S. Spatt, 2015, "Equity trading in the 21st century: An update," *Quarterly Journal of Finance*, 5(1), 1–55.

Aquilina, M., E. B. Budish, and P. O'Neill, 2021, "Quantifying the high-frequency trading 'arms race': A simple new methodology and estimates," *Quarterly Journal of Economics*, forthcoming.

Baron, M., J. Brogaard, B. Hagströmer, and A. Kirilenko, 2019, "Risk and return in high-frequency trading," *Journal of Financial and Quantitative Analysis*, 54(3), 993–1024.

Biais, B., T. Foucault, and S. Moinas, 2015, "Equilibrium fast trading," *Journal of Financial Economics*, 116(2), 292–313.

Boehmer, E., D. Li, and G. Saar, 2018, "The competitive landscape of high-frequency trading firms," *Review of Financial Studies*, 31(6), 2227–2276.

Brogaard, J., A. Carrion, T. Moyaert, R. Riordan, A. Shkilko, and K. Sokolov, 2018, "High frequency trading and extreme price movements," *Journal of Financial Economics*, 128(2), 253–265.

Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan, 2015, "Trading fast and slow: Colocation and liquidity," *Review of Financial Studies*, 28, 3407–3443.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, "High-frequency trading and price discovery," *Review of Financial Studies*, 27(8), 2267–2306.

Budish, E., P. Cramton, and J. Shim, 2015, "The high-frequency trading arms race: Frequent batch auctions as a market design response," *Quarterly Journal of Economics*, 130(4), 1547–1621.

Cespa, G., and X. Vives, 2022, "Exchange competition, Entry, and Welfare," *Review of Financial Studies*, forthcoming.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2014, "Rise of the machines: Algorithmic trading in the foreign exchange market," *Journal of Finance*, 69(5), 2045–2084.

Chakrabarty, B., P. K. Jain, A. Shkilko, and K. Sokolov, 2021, "Unfiltered Market Access and Liquidity: Evidence from the SEC Rule 15c3-5," *Management Science*, 67, 1183–1198.

Chakrabarty, B., R. Pascual, and A. Shkilko, 2015, "Evaluating trade classification algorithms: Bulk volume classification versus the tick rule and the Lee-Ready algorithm," *Journal of Financial Markets*, 25, 52–79.

Chordia, T., T. C. Green, and B. Kottimukkalur, 2018, "Rent seeking by low-latency traders: Evidence from trading on macroeconomic announcements," *Review of Financial Studies*, 31(12), 4650–4687.

Comerton-Forde, C., V. Grégoire, and Z. Zhong, 2019, "Inverted fee structures, tick size, and market quality," *Journal of Financial Economics*, 134(1), 141–164.

Comerton-Forde, C., and T. J. Putniņš, 2015, "Dark trading and price discovery," *Journal of Financial Economics*, 118(1), 70–92.

Conrad, J., S. Wahal, and J. Xiang, 2015, "High-frequency quoting, trading, and the efficiency of prices," *Journal of Financial Economics*, 116(2), 271–291.

Conrad, J. S., K. M. Johnson, and S. Wahal, 2001, "Institutional trading and soft dollars," *Journal of Finance*, 56(1), 397–416.

Corwin, S. A., and P. Schultz, 2012, "A simple way to estimate bid-ask spreads from daily high and low prices," *Journal of Finance*, 67(2), 719–760.

Eaton, G. W., P. J. Irvine, and T. Liu, 2021, "Measuring institutional trading costs and the implications for finance research: The case of tick size reductions," *Journal of Financial Economics*, 139(3), 832–851.

Foley, S., A. Liu, K. Malinova, A. Park, and A. Shkilko, 2020, "Cross-subsidizing liquidity," *Working Paper*, Macquarie University.

Foucault, T., R. Kozhan, and W. W. Tham, 2017, "Toxic arbitrage," *Review of Financial Studies*, 30(4), 1053–1094.

Goldstein, M. A., and K. A. Kavajecz, 2000, "Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE," *Journal of Financial Economics*, 56(1), 125–149.

Haslag, P. H., and M. C. Ringgenberg, 2020, "The demise of the NYSE and NASDAQ: Market quality in the age of market fragmentation," *Working Paper*, Vanderbilt University.

Hendershott, T., and C. M. Jones, 2005, "Island goes dark: Transparency, fragmentation, and regulation," *Review of Financial Studies*, 18(3), 743–793.

Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, "Does algorithmic trading improve liquidity?," *Journal of Finance*, 66, 1–33.

Hendershott, T., and P. C. Moulton, 2011, "Automation, speed, and stock market quality: The NYSE's hybrid," *Journal of Financial Markets*, 14(4), 568–604.

Hoffmann, P., 2014, “A dynamic limit order market with fast and slow traders,” *Journal of Financial Economics*, 113(1), 156–169.

Hou, K., and T. J. Moskowitz, 2005, “Market frictions, price delay, and the cross-section of expected returns,” *Review of Financial Studies*, 18(3), 981–1020.

Indriawan, I., R. Pascual, and A. Shkilko, 2021, “On the Effects of Continuous Trading,” *Working Paper*, Auckland Institute of Technology.

Jones, C. M., and M. L. Lipson, 2001, “Sixteenths: direct evidence on institutional execution costs,” *Journal of Financial Economics*, 59(2), 253–278.

Jovanovic, B., and A. J. Menkveld, 2016, “Middlemen in limit order markets,” Working paper, New York University.

Korajczyk, R. A., and D. Murphy, 2019, “High-frequency market making to large institutional trades,” *Review of Financial Studies*, 32(3), 1034–1067.

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

Lo, A. W., and A. C. MacKinlay, 1988, “Stock market prices do not follow random walks: Evidence from a simple specification test,” *Review of Financial Studies*, 1(1), 41–66.

Menkveld, A. J., and M. A. Zoican, 2017, “Need for speed? Exchange latency and liquidity,” *Review of Financial Studies*, 30(4), 1188–1228.

Pagnotta, E. S., and T. Philippon, 2018, “Competing on speed,” *Econometrica*, 86(3), 1067–1115.

Puckett, A., and X. Yan, 2011, “The interim trading skills of institutional investors,” *Journal of Finance*, 66(2), 601–633.

Roşu, I., 2019, "Liquidity and Information in Limit Order Markets," *Journal of Financial and Quantitative Analysis*, pp. 1–48.

Shkilko, A., and K. Sokolov, 2020, "Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs," *Journal of Finance*, 75(6), 2899–2927.

Spatt, C. S., 2020, "Is equity market exchange structure anti-competitive," *Working Paper*, Carnegie Mellon University.

van Kervel, V., and A. J. Menkveld, 2019, "High-Frequency Trading around Large Institutional Orders," *Journal of Finance*, 74(3), 1091–1137.

Zhu, H., 2014, "Do dark pools harm price discovery?," *Review of Financial Studies*, 27(3), 747–789.

Table 1: Descriptive Statistics

The table reports sample summary statistics. In Panel A, these include market capitalization, daily trading volume, share price, and price volatility, with the latter computed as the difference between the day's high and low prices scaled by the high price. Panel B reports liquidity metrics computed from TAQ, including quoted, effective, and realized spreads as well as price impacts. Panel C reports Abel Noser institutional daily trading volume and execution shortfall. The TAQ data cover the entire sample period of 2003-2017, whereas Abel Noser data end in 2013. Panel D reports price efficiency statistics including the variance ratio, return autocorrelation, and price delay computed at various horizons. We first compute averages for each stock during the sample period and then compute the reported statistics across stocks.

	Mean	St. dev.	25^{th}	Median	75^{th}
Panel A: Sample characteristics					
Market capitalization, $B	11.74	30.34	1.27	2.92	8.79
Volume, sh. '000	2,401	6,378	312	773	2,297
Price, $	47.00	77.46	24.83	38.04	52.90
Volatility	0.03	0.01	0.02	0.03	0.03
Panel B: TAQ liquidity metrics, bps.					
Quoted spread	15.84	16.60	6.27	11.32	19.71
Effective spread	14.38	13.67	6.70	10.49	18.06
Price impact	7.35	6.85	3.43	5.61	9.54
Realized spread	5.57	6.45	2.34	3.62	6.49
Panel C: Abel Noser metrics					
Institutional volume, sh. '000	220	420	40	89	229
Execution shortfall, bps	13.15	9.94	7.13	12.47	18.95
Panel D: Price efficiency metrics					
Variance ratio (15s, 60s)	0.131	0.017	0.121	0.128	0.137
Variance ratio (60s, 300s)	0.204	0.015	0.195	0.202	0.210
Return autocorrelation (15s)	0.062	0.009	0.056	0.059	0.064
Return autocorrelation (300s)	0.126	0.007	0.122	0.126	0.131
Price delay (15s)	0.393	0.157	0.284	0.406	0.509
Price delay (300s)	0.471	0.108	0.398	0.463	0.532

Table 2: Liquidity Costs: TAQ Data

The table contains univariate results (Panel A) and regression coefficient estimates (Panel B) that measure the effects of exchange innovation on quoted and effective spreads. Panel A compares the metrics in the pre-implementation window to those in the post-implementation window. Panel B reports coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing the quoted or effective spread for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

	Quoted Spread				Effective Spread			
	[1]		[2]		[3]		[4]	
Panel A: Univariate results, bps								
Pre	15.54				13.96			
Post	14.67	***			13.04	***		
Panel B: Regression results								
Post	-0.011	***	-0.008	***	-0.009	***	-0.006	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Volume			-0.234	***			-0.005	***
			(0.00)				(0.00)	
Volatility			0.338	***			0.264	***
			(0.00)				(0.00)	
Intercept	0.028	***	0.019	***	0.022	***	0.015	***
	(0.00)		(0.00)		(0.00)		(0.00)	

Table 3: Liquidity Cost Components: TAQ Data

The table contains univariate results (Panel A) and regression coefficient estimates (Panel B) that measure the effects of exchange innovation on realized spreads and price impacts. Panel A compares the metrics in the pre-implementation window to those in the post-implementation window. Panel B reports coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing one of the spread components (the price impact or realized spread) for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** and ** denote the 0.01 and 0.05 levels of statistical significance.

	Realized Spread				Price Impact			
	[1]		[2]		[3]		[4]	
Panel A: Univariate results, bps								
Pre	5.10				7.67			
Post	4.64	***			7.56	***		
Panel B: Regression results								
Post	-0.003	***	-0.003	**	-0.004	***	-0.001	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Volume			0.045	***			-0.052	***
			(0.00)				(0.00)	
Volatility			0.010	***			0.223	***
			(0.00)				(0.00)	
Intercept	0.008	***	0.008	**	0.009	***	0.004	***
	(0.00)		(0.00)		(0.00)		(0.00)	

Table 4: Difference-in-Differences Analysis

The table contains coefficient estimates from a difference-in-differences regression that compares liquidity metrics for the main sample of U.S. stocks and a sample of matched Canadian stocks. To find suitable matches, we select Canadian stocks that are closest to the sample stocks by market capitalization and yearly trading volume estimated during the previous sample year. Our Canadian data are daily and resemble those available from CRSP. As such, in lieu of liquidity metrics we use the low-frequency effective spread estimators proposed by Corwin and Schultz (2012) and Abdi and Ranaldo (2017) as well as the end-of-day (EOD) quoted spread. The regression setup is similar to that discussed in Table 2, aside from using differences between the U.S. and Canadian variables:

$$\Delta DepVar_{ijt} = \alpha_{ij} + \beta_1 Post_t + \beta_2 \Delta Volume_{ijt} + \beta_3 \Delta Volatility_{ijt} + \varepsilon_{ijt},$$

where $\Delta DepVar_{ijt}$ is the difference between the spread proxies (i.e., EOD, CS, and AR) computed on day t for a pair of stocks that includes a U.S. sample stock i and its matched Canadian counterpart j, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, and $\Delta Volume_{ijt}$ and $\Delta Volatility_{ijt}$ controls are the differences between the volume and volatility estimates for each pair of stocks i and j on each day t. If more than one patent is filed during the implementation period, *Post* equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

	Quoted spread		Effective spread			
	EOD		Corwin-Schultz		Abdi-Ranaldo	
	[1]		[2]		[3]	
Post	-0.012	***	-0.010	***	-0.005	***
	(0.00)		(0.00)		(0.00)	
ΔVolume	-0.117	***	0.096	***	-0.067	***
	(0.00)		(0.00)		(0.00)	
ΔVolatility	0.327	***	0.100	***	0.169	***
	(0.00)		(0.00)		(0.00)	
Intercept	0.024	***	0.015	***	0.004	
	(0.00)		(0.00)		(0.00)	
R^2	0.082		0.027		0.021	

Table 5: Volume and Volatility: TAQ Data

The table contains univariate results (Panel A) and regression coefficient estimates (Panel B) that measure the effects of exchange innovation on trading volume and price volatility. Panel A compares the metrics in the pre-implementation window to those in the post-implementation window. Panel B reports coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing the log share volume or volatility computed as the difference between the day's high and low prices scaled by the average of the two for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. We do not use volume and volatility controls in the regressions for, respectively, volume and volatility. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Panel C reports the β_1 coefficient estimates from the DID model estimated using equation (7). Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

	Volume		Volatility	
	[1]		[2]	
Panel A: Univariate results, bps				
Pre	13.34		0.027	
Post	13.36	***	0.026	***
Panel B: Regression results				
Post	0.006	***	-0.010	***
	(0.00)		(0.00)	
Volume			0.494	***
			(0.00)	
Volatility	0.494	***		
	(0.00)			
Intercept	-0.015	***	0.026	***
	(0.00)		(0.00)	
Panel C: DID results				
Post	0.005	***	-0.004	***
	(0.00)		(0.00)	

Table 6: Institutional Trading Costs and Volume: Abel Noser Data

The table contains univariate results (Panel A) and regression coefficient estimates (Panel B) that measure the effects of exchange innovation on execution shortfall and institutional trading volume. Panel A compares the metrics in the pre-implementation window to those in the post-implementation window. Panel B reports coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing the execution shortfall and log share volume for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is total trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. We do not use the volume control in the regression for institutional volume due to high correlation between the two variables. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

	Shortfall		Volume	
	[1]		[2]	
Panel A: Univariate results, bps				
Pre	9.34		10.48	
Post	10.28	***	10.44	***
Panel B: Regression results				
Post	0.001	***	-0.001	***
	(0.00)		(0.00)	
Volume	0.007	***		
	(0.00)			
Volatility	0.074	***	0.165	***
	(0.00)		(0.00)	
Intercept	-0.003	***	0.003	***
	(0.00)		(0.00)	

Table 7: Price Efficiency: TAQ Data

The table contains univariate results (Panel A) and regression coefficient estimates (Panel B) that measure the effects of exchange innovation on price efficiency. We use three efficiency metrics: the variance ratio, return autocorrelation, and price delay. The variance ratio is computed for two intervals: (i) 15 and 60 seconds and (ii) 60 and 300 seconds. The remaining two metrics are computed for 15-second and 300-second intervals. Panel A compares the metrics in the pre-implementation window to those in the post-implementation window. Panel B reports coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing one of the three price efficiency metrics for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is total trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

	Variance Ratio				Return Autocorrelation				Price Delay			
	(15s, 60s)		(60s, 300s)		15s		300s		15s		300s	
	[1]		[2]		[3]		[4]		[5]		[6]	
Panel A: Univariate results												
Pre	0.132		0.204		0.062		0.127		0.404		0.477	
Post	0.130	***	0.203	***	0.060	***	0.126	***	0.392	***	0.482	***
Panel B: Regression results												
Post	-0.002	***	0.000		-0.002	***	0.000		-0.016	***	-0.004	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	
Volume	0.016	***	0.032	***	0.010	***	0.035	***	0.023	***	0.158	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	
Volatility	0.041	***	-0.007		0.044	***	-0.039		0.013	***	-0.087	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	
Intercept	0.004	***	0.001	***	0.005	***	0.001	***	0.042	***	0.011	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	

Table 8: The Cross-Section: Liquidity, Volume, and Volatility

The table reports coefficient estimates from the following regression model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Post_t \times Medium_{it} + \beta_3 Post_t \times Large_{it} + \beta_4 Volume_{it} + \beta_5 Volatility_{it} + \varepsilon_{it},$$

where *Large* and *Medium* are dummy variables for the large and medium stocks, and all other variables are previously defined. The *Post* dummy absorbs the effect for the small stocks, and the two interaction variables capture the effects for the other two size categories *in addition to* the small stock effect. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** and ** denote the 0.01 and 0.05 levels of statistical significance. We do not report the coefficient estimates for the control variables to conserve space.

	[1]		[2]		[3]		[4]	
	TAQ							
	Quoted Spread		Effective Spread		Realized Spread		Price Impact	
Post	-0.005	***	-0.006	***	-0.003	***	-0.001	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Post × Medium	-0.002	***	0.000		0.000		0.000	
	(0.00)		(0.00)		(0.00)		(0.00)	
Post × Large	-0.005	***	0.000		-0.001	**	0.000	
	(0.00)		(0.00)		(0.00)		(0.00)	

	[5]		[6]		[7]		[8]	
	TAQ				Abel Noser			
	Volume		Volatility		Shortfall		Volume	
Post	0.006	***	-0.009	***	0.001		-0.002	***
	(0.00)		(0.00)		(0.00)		(0.00)	
Post × Medium	-0.001		-0.001		0.001	**	0.001	
	(0.00)		(0.00)		(0.00)		(0.00)	
Post × Large	-0.001		-0.002	***	0.002	***	0.001	
	(0.00)		(0.00)		(0.00)		(0.00)	

Table 9: The Cross-Section: Price Efficiency

The table reports coefficient estimates from the following regression model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Post_t \times Medium_{it} + \beta_3 Post_t \times Large_{it} + \beta_4 Volume_{it} + \beta_5 Volatility_{it} + \varepsilon_{it},$$

where *Large* and *Medium* are dummy variables for the large and medium stocks, and all other variables are previously defined. The *Post* dummy absorbs the effect for the small stocks, and the two interaction variables capture the effects for the other two size categories *in addition to* the small stock effect. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance. We do not report the coefficient estimates for the control variables to conserve space.

	Variance Ratio				Return Autocorrelation				Price Delay			
	(15s, 60s)		(60s, 300s)		15s		300s		15s		300s	
	[1]		[2]		[3]		[4]		[5]		[6]	
Post	-0.001	***	0.000		-0.002	***	0.000		-0.014	***	-0.006	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	
Post × Medium	0.001		0.000		0.001	***	0.000		-0.002	***	0.002	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	
Post × Large	-0.002	***	-0.001		-0.002	***	0.000		-0.004	***	0.003	***
	(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)	

Table 10: Robustness

The table reports the results of robustness tests that estimate equation (6) regression coefficients using a set of adjustments to the length of the sample period and to the structure of event windows. The first line reports the base case, which we use throughout the paper. The base case covers the 2003-2017 sample period and uses the 3-3-3 event window, that is a 3-month pre-implementation period, a 3-month implementation period, and a 3-month post-implementation period. The second line examines the sample period from 2003 through 2013, for which we have Abel Noser data. The third line reverts to the 2003-17 sample period, but omits 2008 and 2009 (the financial crisis years). The remaining three lines examine adjustments to the event window, shrinking the implementation period from three months to two and then to one (lines four and five) and then expanding the pre- and post-implementation periods to four and five months (lines six and seven). The coefficient estimates are from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing the quoted, effective, or realized spread, price impact, log share volume, volatility, execution shortfall, and institutional volume computed for each stock i on day t, and all other variables are defined previously. All variables are detrended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Given the size of the table, we omit the asterisks. Instead, we note that the coefficient estimates reported in regular font are statistically significant at the 0.01 level, and the coefficient reported in italicized font – at the 0.05 level. There are no coefficient estimates with statistical significance above the 0.05 level.

	TAQ						Abel Noser	
	QS	ES	RS	PI	Vol	Volat	ExShortf	InstVol
	[1]	[2]	[3]	[4]	[5]	[6]	[7]	[8]
1. 2003-17; 3-3-3	-0.008	-0.006	*-0.003*	-0.001	0.006	-0.010		
(base case)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)		
2. 2003-13; 3-3-3	-0.019	-0.012	-0.005	-0.005	0.007	-0.025	0.001	-0.001
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
3. w/o 2008-09; 3-3-3	-0.006	-0.006	-0.003	-0.002	0.005	-0.008	0.001	*-0.002*
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
4. 2003-17; 3-2-3	-0.025	-0.012	-0.006	-0.001	0.006	-0.018	0.001	-0.005
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
5. 2003-17; 3-1-3	-0.020	-0.010	-0.004	-0.004	0.004	-0.016	0.001	-0.006
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
6. 2003-17; 4-1-4	-0.013	-0.008	-0.003	-0.002	0.005	-0.010	0.001	-0.003
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
7. 2003-17; 5-1-5	-0.005	-0.004	-0.001	-0.001	0.002	-0.007	0.001	-0.002
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)

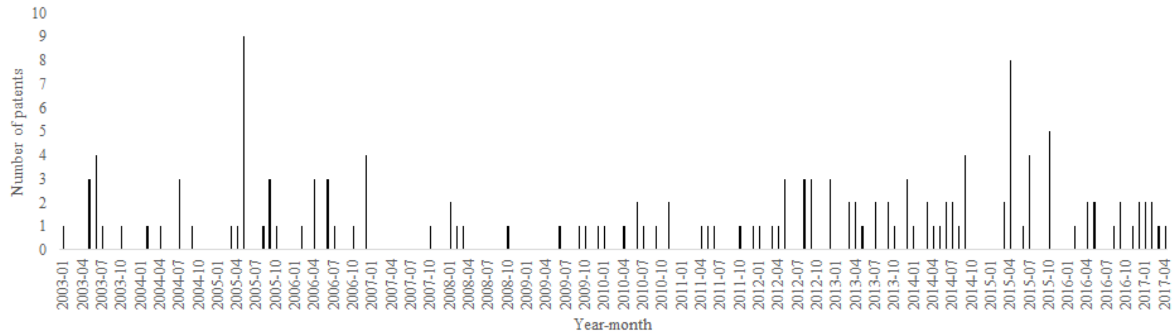

Figure 1. Time series of patent filings
The figure plots monthly patent filings related to equity trading by U.S. exchanges.

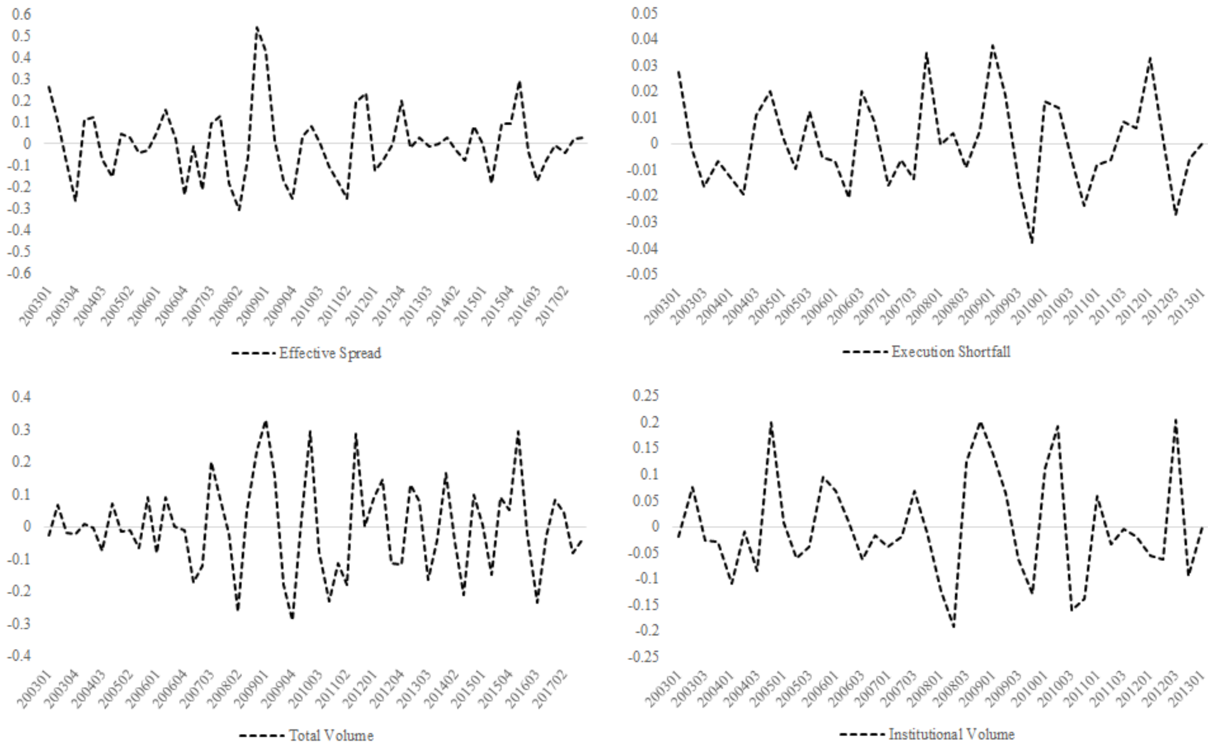

Figure 2. Examples of detrended variables

The figure plots the residuals from regressions of effective spread, execution shortfall, total volume, and institutional volume for each sample stock on a time trend. We perform the same trend adjustments for all variables of interest.

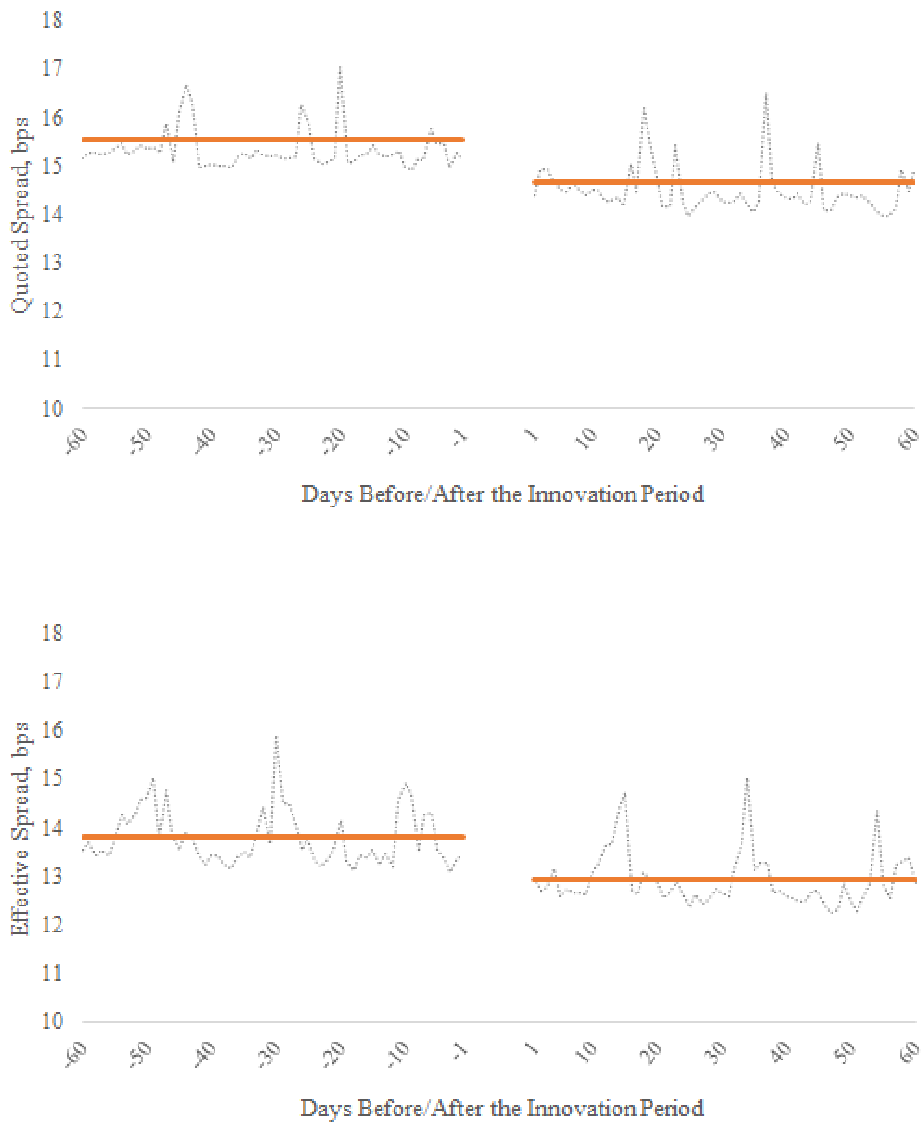

Figure 3. Displayed liquidity and trading costs around patent filings

The figures report quoted and effective spread estimates (dotted lines) and their averages (solid orange lines) during the pre- and post-innovation periods.

Appendix to “Who Benefits When Exchanges Innovate?”

A.1 Sample statistics estimated across time

In Table 1 of the main text, we report summary statistics estimated across stocks. For instance, the standard deviation of 16.60 reported for quoted spreads measures the dispersion of stock-level means. While informative about the sample, this statistic is less useful when interpreting coefficient estimates from fixed effect panel regressions that we use throughout the paper. Therefore, when discussing economic significance we rely on the time-series summary statistics, and more specifically the standard deviations estimated across sample days rather than across sample stocks. For comparison to the above-mentioned figure, the standard deviation for quoted spreads is 6.20 when computed via this method. We report these standard deviations, along with the other summary statistics in Table A1.

A.2 Additional price efficiency results

In Table 7 of the main manuscript, we report price efficiency statistics for two select estimation horizons. In Table A2, we augment these results with additional horizons. Specifically, we report the variance ratios estimated at (10s, 60s), (15s, 60s), (30s, 60s), and (60s, 300s) horizons, and we report return autocorrelations and the price delay metrics computed at 10, 15, 30, 60, and 300s horizons. As we mention in the main manuscript, price efficiency appears to improve post-innovation, and the improvements are most prominent at shorter estimation horizons.

Table A1: Descriptive Statistics: Time-Series

The table reports sample summary statistics computed across sample days. In comparison, Table 1 in the main text computes summary statistics across sample stocks. The statistics include liquidity metrics computed from TAQ, including quoted, effective, and realized spreads as well as price impacts. We also report Abel Noser institutional daily trading volume and execution shortfall. Finally, we report price efficiency statistics (i.e., the variance ratio, return autocorrelation, and the price delay) estimated at the same horizons as those used in the main manuscript. The TAQ data cover the entire sample period of 2003-2017, whereas Abel Noser data end in 2013. We first compute averages for each day during the sample period and then compute the reported statistics across days.

	Mean	St. dev.	25th	Median	75th
Volume, sh. '000	2,461	1,191	1,660	2,177	2,945
Volatility	0.03	0.01	0.02	0.03	0.03
Quoted spread, bps	15.93	6.20	12.20	14.87	18.47
Effective spread, bps	14.39	8.04	9.71	12.37	16.50
Price impact, bps	7.38	6.62	3.59	6.30	9.88
Realized spread, bps	5.61	8.34	1.18	4.65	8.82
Institutional volume, sh. '000	218	232	68	141	280
Execution shortfall, bps	13.03	158.85	-57.40	6.88	79.26
Variance ratio (15s, 60s)	0.131	0.100	0.052	0.110	0.188
Variance ratio (60s, 300s)	0.204	0.149	0.084	0.176	0.294
Return autocorrelation (15s)	0.061	0.046	0.025	0.052	0.088
Return autocorrelation (300s)	0.126	0.094	0.050	0.107	0.183
Price delay (15s)	0.393	0.150	0.288	0.381	0.484
Price delay (300s)	0.468	0.240	0.304	0.463	0.635

Table A2: Additional price efficiency results

The table contains regression coefficient estimates that measure the effects of exchange innovation on price efficiency. We use three efficiency metrics: the variance ratio, return autocorrelation, and price delay. The variance ratio is computed for four intervals: (i) 10 and 60 seconds, (ii) 15 and 60 seconds, (iii) 30 and 60 seconds, and (iv) 60 and 300 seconds. The remaining two metrics are computed at 10, 15, 30, 60, and 300-second intervals. We report coefficient estimates from the following model:

$$DepVar_{it} = \alpha_i + \beta_1 Post_t + \beta_2 Volume_{it} + \beta_3 Volatility_{it} + \varepsilon_{it},$$

where $DepVar_{it}$ is the estimated residual from equation (5) representing one of the three price efficiency metrics for each stock i on day t, $Post_t$ is a dummy variable equal to one during the post-innovation period and zero during the pre-innovation period, $Volume_{it}$ is total trading volume, and $Volatility_{it}$ is the difference between the high and low prices on day t scaled by the high price. If more than one patent is filed during the implementation period, $Post_t$ equals to the number of patents. This approach allows us to identify the effects per patent. All variables are de-trended and standardized, and as such the models control for stock fixed effects. Standard errors in parentheses are double-clustered across firms and over time. Asterisks *** denote the 0.01 level of statistical significance.

Panel A: Variance ratio					
	(10s, 60s)	(15s, 60s)	(30s, 60s)	(60s, 300s)	
	-0.002 ***	-0.002 ***	-0.001 ***	0.000	
	(0.00)	(0.00)	(0.00)	(0.00)	
Panel B: Return autocorrelation					
	10s	15s	30s	60s	300s
	-0.002 ***	-0.002 ***	-0.002 ***	-0.001 ***	0.000
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Panel C: Price delay					
	10s	15s	30s	60s	300s
	-0.017 ***	-0.016 ***	-0.014 ***	-0.011 ***	-0.004 ***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)