Review of Finance - preprint, 2022, 1–42
doi:

Tick Size Wars: The Market Quality Effects of Pricing Grid Competition

Sean Foley,[1,*] Tom G. Meling[2] and Bernt Arne Ødegaard[3]

[1]Macquarie University, [2]Fisher College of Business, The Ohio State University and [3]University of Stavanger (UiS)

[*]Sean Foley sean.foley@mq.edu.au

FOR PUBLISHER ONLY Received on 0 October 2020; revised on 0 February 2022

Abstract

We explore the effects of a "tick size war" in which European trading venues directly competed on the minimum pricing increment in the limit order book, the tick size. We find that venues that reduced their tick size immediately captured market shares of both quoted and executed volume from the exchanges that kept their ticks large. We find that tick size competition improves market quality, reducing trading costs and increasing market-wide depth and volume. These market quality improvements are strongest in stocks where the bid-ask spread was constrained to one tick, where liquidity providers use the finer pricing grid to engage in price competition.

Key words: Equity Trading; Limit Order Markets; Tick Sizes

1. Introduction

Fragmentation is a key feature of modern equity markets. Spurred by regulatory initiatives introducing competition to monopolistic primary exchanges, notably Regulation NMS in the United States and MiFID in Europe, a plethora of new trading venues have emerged, evolving various tools to attract order flow. With the new regulations came a regulator-mandated minimum tick size, ensuring that exchanges could not compete by changing the granularity of their pricing schedules. Central to the ensuing competition for order flow has been the desire to circumvent these regulations and "synthetically" compete on tick sizes through alternative trading mechanisms such as dark pools, batch auctions, and inverted maker-taker price structures.

In this paper, we examine the market quality effects of *explicit* tick size competition between trading venues. Our laboratory is the European "tick size war" of 2009, in which owing to the absence of a regulator-mandated uniform tick size, entrant trading venues were able to undercut the tick sizes of the primary exchanges. We find that competitive tick size reductions lower market-wide trading costs, improve liquidity and informational efficiency, and drive both quoted and executed market share away from large-tic trading venues, towards the small-tick venues.

Empirically, the connection between equity market fragmentation and a desire to circumvent existing tick size regimes has been documented in the context of inverted maker-taker venues in recent papers by for example Spatt (2017) and Comerton-Forde et al. (2019). The proliferation of dark pool venues in the United States is shown by Kwan et al. (2015) to be primarily motivated by the desire to navigate around tick sizes. Such trading venues typically enable liquidity providers to "undercut" the lit limit order book by an economically insignificant amount, gaining time priority without significant costs. The undesirable nature of these practices has led regulators in countries such as Australia and Canada to introduce regulations that require "meaningful" price improvement (Foley and Putniņš, 2016).[1]

Theoretically, the incentives for profit-motivated, competitive trading venues to attempt to subvert existing tick size regimes are highlighted by Chao et al. (2018). Their model argues that "competition on fees and fragmentation among exchanges is a type of financial innovation in response to mandated discrete tick sizes." This argument is consistent with Angel et al. (2011) who argue that sub-tick trading helps electronic liquidity providers (ELPs) gain queue priority in the order book. Foucault et al. (2013) argue that discrete tick sizes may prevent prices from neutralizing the impact of fees and rebates, providing further incentives for competitive venues to subvert them. These theoretical tensions make assessing the role of tick size competition on market quality and fragmentation an empirical question.

We study the effects of the European "tick size wars" of 2009, in which trading venues competed directly on the granularity of their pricing grids to gain market share—one of very few such events in modern times. The tick size war raged for several months until an accord led to the harmonization of pricing schedules across all competing venues. A major advantage of our empirical setting is that we directly observe the results of tick size competition (and its removal), allowing us to understand how both liquidity suppliers and market participants respond to the evolving landscape, without having to

[1] Most recently, the US Securities and Exchange Commission's (SEC's) tick size pilot examined the role tick sizes play in encouraging both liquidity and new listings while simultaneously balancing the desire for reduced transactions costs. Optimal tick size regimes are typically examined through the lens of market-wide changes (see e.g. Rindi and Werner (2019) for a discussion of the US tick size pilot) or through more firm-specific situations such as stock splits and re-consolidations (Yao and Ye, 2018).

deal with the distortions inherent in alternative mechanisms, such as inverted maker-taker fees or dark trading. This allows us to isolate the effects of exchange competition based solely on pricing grids.

Our main finding is that the trading venues that reduced their minimum price increments during the tick size war immediately captured market shares of both quoted and executed volume from the primary exchanges that kept their price increments wide, with liquidity suppliers particularly attracted to the small-tick venues. The existing liquidity providers improved the prices possible on the large-tick exchanges, often competing with each other to quote at the best prices on the small-tick venues. We find that this competition improved aggregate market quality — reducing transaction costs, increasing order book depth, and improving the efficiency of market-wide prices. We find that these improvements to aggregate market quality are strongest in stocks where the bid-ask spread was constrained by the old tick size. Our findings are consistent with the notion that the "one size fits all" approach currently taken to tick size regulation globally may require revision.

We make three main contributions to the literature. The first is to show that the ability to directly undercut the tick sizes of competing exchanges is an important driver of exchange volume, which by itself can generate fragmentation. This is consistent with the theoretical arguments of Colliard and Foucault (2012) that exchanges will compete towards continuous prices to neutralize frictions. Given discrete prices, Foucault et al. (2013) predict the existence of competition based on fees or, in our context, the pricing grid itself. Chen and Duffie (2021) argue that fragmentation can exist in equilibrium owing to the reduced price impact and increased informational efficiency created by inter-market competition. We observe evidence of such an equilibria during retaliation events, in which the primary exchanges competitively respond with reduced tick sizes. Such behavior, we find, leads to a significant increase in quote matching on the incumbent venue, but does not significantly reverse the increased fragmentation, with traded volume remaining split between the incumbent and entrant venues.

This confirmation of the importance of tick size competition is also consistent with recent work by Chao et al. (2018) arguing that exchanges will compete to get as close as possible to a continuous pricing grid, differentiating themselves on cum-fee pricing. These results, which are directly observed in our analysis via tick size competition, further connect to the empirical findings of studies such as Kwan et al. (2015), Foley and Putniņš (2016), and Menkveld et al. (2017) on the desirability of dark pools, as well as Menkveld (2013), Battalio et al. (2016) and Comerton-Forde et al. (2019) on why fee structures (in general) and inverted maker-taker venues (in particular) may be preferred to traditional venues. In our setting, traders are observed to utilize the option to undercut quotes by a very small fraction to attract order flow, consistent with dark pricing "between" the bid-ask spread, and the sub-tick levels of price improvement provided by inverted maker-taker venues. Such undercutting is most prominent, we observe, in securities that are heavily constrained by the tick size. Showing that these effects persist when trading venues are allowed to directly compete on the pricing grid suggests that many of the indirect competitive mechanisms may themselves be driven by the desire of market participants for a finer pricing grid.

Our second main contribution is to show that the market quality gains from finer pricing grids are concentrated in the securities for which bid-ask spreads are most constrained by the tick size. We utilize a measure of "optimal" tick sizes developed by Kwan et al. (2015) to show that when spreads are heavily constrained, the introduction of a new, smaller tick size results in market participants driving down transactions costs and increasing quoted depth. This occurs because liquidity providers compete on the finer pricing grid, narrowing the spread by several new ticks. These results are consistent with the theoretical findings of Anshuman and Kalay (1998) that discrete pricing grids generate rents for market makers and discourage information production. Our results are also in line with the findings of Goettler et al. (2005) that smaller tick sizes reduce transactions costs, as liquidity providers are able to quote "closer to the consensus value." The results are also consistent with the empirical findings of Bessembinder (2003) and Chakravarty et al. (2004) that decimalization in the United States primarily benefited the most liquid securities, as well as the findings of Kerr et al. (2020) and Albuquerque et al. (2020) that market efficiency improves with smaller tick sizes.

Our third main contribution is to document the effects of tick size competition on trader behavior, particularly for ELPs. We find evidence of two different competitive scenarios, one in which the ELPs "undercut" the existing pricing grid by one small tick, leaving the primary quote setting functions at the incumbent exchange, and another where the narrower pricing grid is aggressively used by market makers to generate new competitive quotes on the entrant exchange which improves on the incumbent exchange by many new ticks. When the ELPs are provided with a finer pricing grid, they predominantly undercut existing prices by the minimum amount possible. Competition among the ELPs forces bid-ask spreads to decrease in the stocks that were previously the most heavily constrained, consistent with the theoretical predictions of Kadan (2006) that more aggressive competition occurs between market makers in the presence of reduced tick sizes. This indicates that identifying and reducing tick sizes for such securities can greatly improve the measures of market quality.

Our findings also contribute to the existing policy debate on optimal tick size regimes. Our results suggest that a new tick size regime—especially for heavily constrained stocks—may halt the continuing innovations in exchange market design, which Chao et al. (2018) argue are focused on subverting market-wide, mandated tick sizes. Such a regulatory change may also help improve market quality across exchanges. Moreover, recent analyses of the SEC's now-concluded tick size pilot provide conflicting evidence on the market quality effects of tick size changes.[2] The inability of this literature to reach a consensus on tick size regimes, combined with the intentions of MiFiD II to reform tick sizes,[3] speaks to the importance of additional evidence, such as ours, for regulators regarding how changes to tick size regimes impact trader behavior and welfare, as well as the efficiency of prices.

More broadly, our findings connect to policy debates on optimal maker-taker regimes, as exemplified by the SEC's and Canadian Securities Administration's plans for pilots to remove maker-taker fees (see, e.g., Malinova et al. 2019) from equity markets altogether. Such an approach seeks to eliminate the most prominent form of cum-fee price competition between trading venues. Based on our findings, it could be prudent for policymakers to consider the addition of a "direct" competition model, particularly among the most tick-constrained stocks, where venues prohibited from adopting potentially distortionary (see, e.g., Battalio et al. 2016) variations on maker-taker models to undertake de-facto price competition are instead allowed to compete on the granularity of their pricing grids.

The paper proceeds as follows. Section 2 provides institutional details and summarizes the main events of tick size war. Section 3 provides details on the data used in the empirical analysis. Section 4 presents our main empirical results. Section 5 analyzes changes to order submission strategies during the tick size war. Finally, Section 6 concludes.

2. Institutional Details

By the early 2000's, the European national stock exchanges, such as those in London, Copenhagen, Oslo and Stockholm, had all completed the transition to pure electronic limit order markets. At the time, the national exchanges held dominant positions in their respective markets. This dominance was broken by the MiFID reform in 2007. MiFID

[2] Chung et al. (2020) show that increased tick sizes have resulted in costs for small orders but benefits for larger orders, owing to reduced trading costs. This contradicts the findings of Griffith and Roseman (2019) of reduced liquidity and resiliency, resulting in reduced asset prices (Albuquerque et al., 2020). Any observed benefits of the tick size pilot have currently been primarily ascribed to HFT participants (Bartlett and McCrary, 2020; O'Hara et al., 2018), as opposed to increased overall liquidity provision, as hoped.

[3] MiFiD II required EU member states to introduce tick sizes based on multiple facets of stocks liquidity by March 2020. In most cases, this resulted in a narrowing of tick sizes.

unleashed competition for European order flow, causing equity trading to fragment across a wide range of trading venues.

2.1. The role of MiFID in fragmenting European equity markets

MiFID introduced a new type of trading venue—the multilateral trading facility (MTF). The first MTFs introduced in Europe were Chi-X in 2007, followed by BATS Europe and Turquoise in 2008.[4] The MTFs initially used similar trading terms as the national exchanges; in particular, they used the same tick size schedules. The MTFs started by trading the largest companies on the national exchanges, and gradually expanded their coverage. By 2009, the MTFs typically offered trading in all the blue-chip index stocks at the national exchanges—that is, the most liquid stocks—whereas their coverage of less-liquid stocks was less comprehensive.

In the summer of 2009, the three largest MTFs, Chi-X, BATS, and Turquoise, unexpectedly lowered their tick sizes for stocks listed on the London Stock Exchange (LSE), Copenhagen Stock Exchange (COP), Oslo Stock Exchange (OSE), and Stockholm Stock Exchange (STO), in an event swiftly referred to as the "tick size war." The MTFs maintained the smaller tick sizes for up to six months, before the exchanges reverted to common tick size schedules. These events are the subject of our analysis.[5]

2.2. The events of the tick size war

The sequence of events is summarized in Table 1. The tick size war can be divided into three phases. In the first phase, which we call the *break-out phase*, the main focus of our empirical analysis, the MTFs Chi-X, Turquoise, and BATS challenged the market positions of the Scandinavian primary exchanges—Copenhagen, Oslo, and Stockholm—by reducing the tick size for their (respective) selections of Danish, Norwegian, and Swedish stocks. The tick size war started on June 1, 2009, when Chi-X reduced its tick sizes for all stocks with a Copenhagen, Oslo, or Stockholm primary listing. Turquoise followed on June 8 by reducing its tick sizes for all Scandinavian stocks and five LSE

[4] For more details on the global roll-out of Chi-X, see He et al. (2015) and Malceniece et al. (2019).

[5] The Scandinavian exchanges have, in recent years, been the testing ground of several empirical studies. For example, using Swedish data, Brogaard et al. (2015) explore the effects of a colocation reform, whereas van Kervel and Menkveld (2019) examines the impact of HFT on institutional trading costs. Using Norwegian data, Næs et al. (2011) explore the connection between stock market liquidity and the business cycle, Skjeltorp and Ødegaard (2015) looks at designated market makers, Jørgensen et al. (2018) investigate order-to-trade ratios, and Meling (2021) explores the effects of post-trade anonymous trading, a topic also studied by Dennis and Sandås (2020). Other recent studies using data on Nordic stock trading include Hvide et al. (2021).

listings. Finally, on June 15, BATS reduced its tick sizes for all Scandinavian stocks, as well as 10 stocks on the LSE and five listed in Milan (BATS, 2009).

Table 1. Timeline—Main events of the tick size war

Time	**Market(s)**	**Event**
2007	Europe	Chi-X established as first MTF post-MiFID.
2008	Europe	BATS and Turquoise established as MTFs.
2008	Sweden	Chi-X enters, smaller tick sizes than STO.
2008	Denmark/Norway	Chi-X enters, same tick sizes as COP/OSE.
June 1, 2009	Scandinavia	Chi-X reduces tick sizes for all stocks.
June 8, 2009	Scandinavia	Turquoise reduces tick sizes for all stocks.
June 8, 2009	UK	Turquoise reduces tick sizes for 5 liquid stocks.
June 15, 2009	Scandinavia	BATS reduces tick sizes for all stocks.
June 15, 2009	UK	BATS reduces tick sizes for 10 liquid stocks.
June 16, 2009	UK	Turquoise reduces tick sizes for 5 liquid stocks.
June 22, 2009	UK	LSE and Chi-X reduce tick sizes for liquid stocks.
July 6, 2009	Norway	OSE retaliates, reduces tick sizes for 25 liquid stocks.
August 26, 2009	Sweden	STO retaliates, reduces tick sizes for 10 liquid stocks.
August 31, 2009	Norway	Harmonization to common tick size schedule.
October 26, 2009	Sweden	Harmonization to common tick size schedule.
January 4, 2010	Denmark	Harmonization to common tick size schedule.

2.3. Tick sizes: The battleground

The MTF tick size reductions during the *break-out phase* were substantial. For example, at the time of Chi-X's June 1 tick size reduction, OSE stocks traded across three tick size schedules, with tick sizes varying between NOK 0.01 and 1, with price cutoffs varying depending on a combination of index membership and stock liquidity. The new Chi-X tick size schedule, by contrast, introduced a NOK 0.001 tick size for all OSE shares (regardless of their index status) with prices below NOK 10, and a NOK 0.005 tick size for shares priced above NOK 10. The tick size schedules introduced by BATS and Turquoise were less aggressive, but still offered substantially smaller tick sizes than the OSE. Similar between-venue tick size differences opened for the trading in Swedish and Danish stocks.[6]

In the second phase of the tick size war—the *retaliation phase*—the primary exchanges responded in kind to the tick size reductions of their competitors. The LSE moved first, with both the LSE and Chi-X matching the smaller tick sizes of BATS and Turquoise after just two weeks (on June 22). On July 6, 2009, about a month into the tick size war, the OSE uniformly reduced its tick size to 0.01 for trading in their 25

[6] Section A of our Internet Appendix provides more detailed summaries of the tick size changes.

most liquid stocks.[7] In doing so, the OSE largely mitigated the between-venue tick size differences created for its most liquid stocks during the *break-out phase*. On August 26, the STO also retaliated, although in a more limited fashion, by reducing the tick size for 10 highly liquid stocks but allowing between-venue tick size differences to persist for less liquid stocks.

2.4. Peace and tick size harmony

The final stage of the tick size war was the *harmonization phase*. On June 30, 2009, the Federation of European Securities Exchanges (FESE) brokered a deal to harmonize tick sizes across the primary exchanges and the MTFs. FESE argued that the recent tick size reductions were not in the interest of investors and that too fine pricing grids could have detrimental effects on market quality. The FESE agreement lead to a pan-European harmonization of tick size schedules, which both simplified and reduced the number of different tick size schedules used by the various exchanges. These changes were to be implemented within six months. The Scandinavian markets responded sequentially; Oslo harmonized its tick sizes first (on August 31, 2009), followed by Stockholm (October 26, 2009), and finally Copenhagen (January 4, 2010). On average, tick sizes increased in Norway for both the OSE and Chi-X following harmonization, whereas in Sweden and Denmark, the primary exchange tick size decreased, but those of the competing venues increased.[8]

3. Data

This section details our data sources, provides an overview of the variables used in the empirical analysis, and presents descriptive statistics.

3.1. Data sources

We obtain trade-and-quote data from the Thomson Reuters Tick History (TRTH) database. For trading venues with displayed limit order books, TRTH provides data on all executed trades, as well as the 10 best levels of the order book. We collect data for the trading in all Scandinavian securities on their primary exchanges (the STO,

---

[7] In a press release, the OSE stated that other venues "offer trading with tick sizes that are significantly lower than Oslo Børs offers. Oslo Børs has therefore found it necessary to respond to these changes." Press release 29/06/2009: Oslo Børs changes the tick size for OBX shares.

[8] Note that in the London market, the tick size war lasted less than two weeks. Our main analysis therefore focuses on the longer-lived events in the Scandinavian markets. Section E of the Internet Appendix provides additional empirical analyses for London-listed securities.

OSE, and COP) as well as for the trading in the same securities on Chi-X, BATS, and Turquoise.[9] In addition, we collect data on the number of shares outstanding and constituency in each market's blue-chip index for the period 2008–2010.

We use the TRTH trade-and-quote data to construct several standard measures of market quality, including the quoted bid-ask spread, effective and realized spreads, the price impact, and measures of order book depth and realized volatility. The construction of these market quality measures follows standard procedures in the market microstructure literature, and the calculations are conducted at the 1-minute horizon where time is required (that is, for realized spreads and price impacts), following the findings of Conrad and Wahal (2020) for the likely holding periods of liquidity providers in 2010. The exact definitions of each of the market quality measures can be found in Section B of the Internet Appendix.

3.2. Measuring aggregate market quality

In addition to the standard venue-specific measures of market quality, we construct measures of aggregate market quality across all the trading venues in a given market. Although the European equity markets have no formal notion of a national best bid and offer (NBBO),[10] we find it useful to construct an estimate of the best bid and offer prices across the competing trading venues.[11] We do so by aligning contemporaneous order books across trading venues and, at each point in time with an update to the order book, recording the best bid and ask prices. For simplicity, we refer to these estimates as the "NBBO." These NBBO prices are then used to construct market-wide measures of transactions costs and realized volatility.

We also construct a measure of market-wide order book depth. This is not straightforward when different venues have different tick sizes. Figure 1 illustrates the issue. By only considering the depth at the best bid and ask on the small-tick venue, one would overlook that the small-tick venue quotes additional depth at prices that still improve on the prices at the large-tick venue. In our view, the most relevant measure is one that aggregates all depth quoted on the small-tick venue at prices equal to (or better than) the best prices at the large-tick venue, as illustrated by the shaded area in

[9] Both the primary exchanges and the MTFs trade securities in their respective native currency. For example, the Norwegian stock Statoil is traded in Norwegian kroner both in Oslo and on the MTFs.

[10] Unlike the American RegNMS, which requires trading at the NBBO, the European regulation only requires "best execution," which examines all features of execution, of which price is one component.

[11] The NBBO is constructed based on all available quotes from the primary exchange (COP, OSE or STO), as well as BATS, Chi-X, and Turquoise.

Figure 1. We view this as the volume at the small-tick venue that competes with the volume available at the best prices on the large-tick venue.

Fig. 1: The across-market depth measure

The figure illustrates the construction of comparable depth measures at exchanges with different tick sizes. The depth at the large-tick exchange is measured as the depth at the best bid and ask on that exchange. On the small-tick exchange, we sum all depth at prices equal to or better than the best bid and ask prices at the large-tick exchange. The shaded areas indicate the comparison.

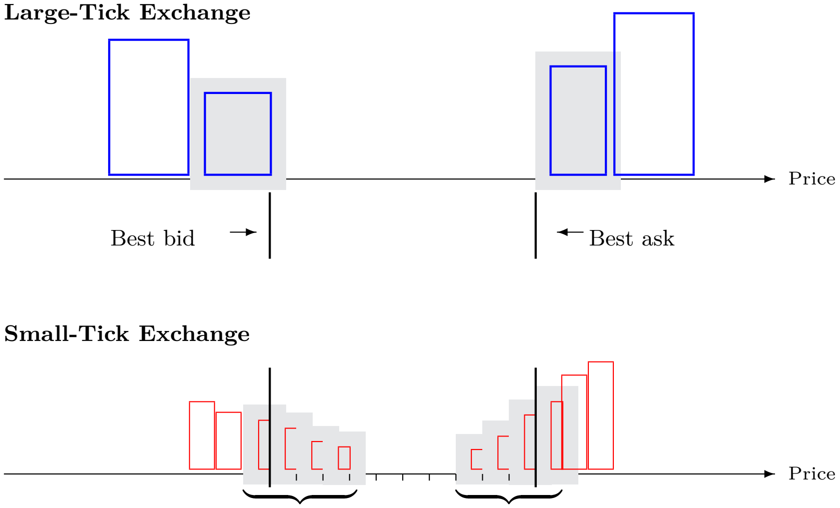

3.3. Descriptive statistics

Panel A of Table 2 provides a count of the number of stocks that, in the period 2008–2010, could be traded on the primary exchanges and the MTFs. Combined, the three primary exchanges offered trading in 1023 stocks, of which 77 were constituents of the exchanges' highly liquid blue-chip indexes. Although their exact stock coverage varied, the MTFs offered trading in a large number of both index and non-index stocks. For example, BATS offered trading in more than 150 Scandinavian stocks, covering significant numbers of index and non-index stocks in each market. Our analysis considers both index and non-index stocks, which allows us to study the effects of the tick size war for both liquid and illiquid stocks.

We compare the market quality in Scandinavian index stocks at the primary exchanges and at the competing MTFs. Table 2 suggests that market quality, as measured by spreads, order book depth, and trading volume, is significantly better at the primary exchanges than the MTFs. For example, the average quoted and effective bid-ask spreads on the MTFs are 1% and 0.32%, respectively, about three times larger than the average on primary exchanges. In terms of depth, the MTFs average about 480 thousand kroner—about a fifth of the primary exchange average of 2.2 million. In terms of executed trading volume, the MTFs average 7.5 million kroner, significantly smaller than the 269 million trading volume on the primary exchanges. Although liquidity seems much better on the primary exchanges, realized volatility is on average considerably smaller on the MTFs—at 0.3% on the MTFs compared with 0.69% on the primary exchanges—perhaps owing to the fact that shares are relatively less frequently traded at the MTFs.

Finally, Panel B of Table 2 summarizes our market-wide spread measures. As constructing an NBBO including the MTFs quotes can only improve the market-wide quoted spread, the average NBBO quoted spread (0.32%), is lower than the 0.36% primary exchange average, but both the effective and realized spreads, which involve trades, are higher for the NBBO than for the primary exchange alone. This is because trades on the MTFs may not necessarily improve on the effective or realized spreads of the primary exchange. Note also that the number of observations differs slightly between the NBBO and MTF samples due to differences in the coverage of stocks.

Table 2. Summary statistics

Panel A provides the number of Scandinavian securities that could be traded at the various trading venues in the period August 2008 to March 2010. For the primary exchanges, COP, OSE, and STO, we group stocks into constituents of the blue-chip indexes, and others. For the competing MTFs, the statistic reported is the total number of traded stocks, regardless of index status. Panel B summarizes our main market quality measures in the period January–May 2009. *Quoted spread*: Difference between the best bid and ask divided by the mid-price, calculated at each update of the order book, and then averaged over the trading day. *Effective spread*, *realized spread* and *price impact* are calculated for each trade, and averaged across all trades during the trading day. The *NBBO* versions of the above four measures use the NBBO to calculate the best bid and ask, instead of the venue-specific best bid and ask. *Realized volatility* is calculated on a daily level as the second (uncentered) moment of 10-minute returns. For each limit order book we calculate *depth* as the sum of the bid and ask depth (in kroner). The daily depth is the average across limit order book states. *Fraction at best quote* is the fraction of the day a given trading venue is quoting the market-wide best price. Finally, *order to trade* is the number of changes to the limit order book divided by the total number of trades that day. In the construction of all market quality measures, we exclude the first and last half hour of the trading day. The volume and market capitalization statistics are reported in local currency (kroner: Denmark–DKK, Norway–NOK, Sweden–SEK). Averages for the primary exchanges (Copenhagen, Oslo and Stockholm) are calculated separately for the stocks inside and outside the blue-chip indexes. Averages for the MTFs only consider index stocks.

Panel A: Equities in sample (2008–2010)

		Stocks in Cross-section				
		Primary Exchange		Chi-X	BATS	Turquoise
Country	Index	Index	Non-index			
Denmark	Copenhagen 20	20	209	29	36	37
Norway	Oslo OBX	26	222	33	46	41
Sweden	Stockholm OMX	30	516	68	97	37

Panel B: Descriptive statistics (January–May 2009)

	Primary Exchanges						MTFs		
	Index stocks			Non-index stocks			Index stocks		
	Mean	Median	N	Mean	Median	N	Mean	Median	N
Quoted spread (%)	0.36	0.33	7686	3.62	2.81	59139	1.02	0.62	13592
Effective spread (%)	0.13	0.12	7685	1.15	0.77	47658	0.32	0.22	12017
Realized spread (%)	0.06	0.05	7685	0.71	0.39	46985	0.13	0.08	11957
Price impact (%)	0.08	0.07	7685	0.48	0.20	46943	0.19	0.11	11956
NBBO quoted spread (%)	0.32	0.27	7686	3.62	2.81	59139	0.26	0.25	11232
NBBO effective spread (%)	0.25	0.19	7685	1.15	0.77	47658	0.27	0.22	14339
NBBO realized spread (%)	0.20	0.13	7685	0.72	0.40	46985	0.24	0.18	14339
NBBO price impact (%)	0.07	0.06	7685	0.48	0.20	46943	0.06	0.05	14339
Realized volatility (%)	0.69	0.41	7686	0.74	0.50	45225	0.30	0.28	12345
Depth (000's)	2240	1016	7686	187	69	59815	488	260	13870
Kroner volume (000,000's)	269.2	130.8	7686	4.1	0.2	64280	7.5	1.3	15294
Fraction at best quote (%)	80.1	95.1	7008	99.8	100.0	50737	35.1	36.1	11232
Market share (%)	96.1	98.0	7686	100.0	100.0	64280	1.9	0.7	15294
Market cap (000,000's)	50748	21440	7686	3242	406	64280			

4. Main Results

We begin our empirical analysis by exploring how the tick size war affected market shares and market quality across the participating trading venues. Digging deeper into the mechanisms underlying the results, we then examine the extent to which a stock's primary exchange bid-ask spreads were constrained by the tick size in the period before the tick size war, which is an important driver of the observed market share and market quality effects. Although the empirical analysis is primarily focused on the *break-out phase* of the tick size war, we also present evidence on the effects of the *retaliation* and *harmonization* phases.

4.1. Did the tick-size-reducing MTFs capture market share?

We consider two measures of market share. The first, a post-trade measure, is simply a trading venue's market share of transacted volume. The second, a pre-trade measure, quantifies how frequently a venue quotes the market-wide best available price.[12] We examine both pre- and post-trade market shares to understand how trading and quoting activity changes in response to between-venue tick size competition.

Figure 2 plots the two market share measures for blue-chip index stocks, separately for the three primary exchanges in our sample: Copenhagen, Oslo, and Stockholm. The graphical evidence is striking. Immediately after the start of the tick size war, on June 1, 2009, we find economically significant reductions in both pre-trade and post-trade primary exchange market shares across all three Scandinavian markets. In terms of executed trading volume, we find market share losses for the primary exchanges that vary between 2.4 percentage points and 3.5 percentage points.[13] Chi-X is the primary benefactor of the primary exchanges' lost market shares, we find, capturing post-trade market shares of between 2.4 and 4 percentage points.

Even more extreme primary exchange market share losses are observed in terms of best quotes. In the Copenhagen and Oslo markets, from the first day of the tick size war, the primary exchanges shift from quoting at the market-wide best price virtually all the time to only about half the time. For Stockholm, the effects are less extreme, possibly reflecting the fact that Chi-X already in the pre-war period offered smaller tick sizes than the STO. Notwithstanding this fact, the observed reduction in pre-trade market share is close to 20 percentage points. For Copenhagen and Oslo, Chi-X is again the

[12] The measure is calculated separately for the bid and ask side and then averaged. Note that this measure does not sum to one, as several exchanges can be (and typically are) quoting the best price.

[13] These numbers represent estimates from a before-after regression specification. The regression estimates are presented in Internet Appendix Section D.

main benefactor of the pre-trade market share losses, whereas for Stockholm, BATS and Turquoise increase their pre-trade market share the most from before to after the start of the tick size war.

The graphical results presented in Figure 2 provide compelling evidence that competing trading venues can capture market shares, in terms of both quoted and executed volume, from incumbent stock exchanges by directly competing on the minimum price increment. Although the primary exchanges' significant loss of market share during June 2009 initially seemed permanent, following each market's tick size harmonization, as Figure 2 shows, the primary exchanges mostly recovered their dominant position in terms of quoting at the market-wide best bid and offer prices, and the gains made by alternative venues in executed volume largely remain. This is consistent with the empirical findings of Foley et al. (2019) that broker connectivity to new venues is "sticky." In Section 4.5, we show that the graphical evidence in Figure 2 is robust to using a more rigorous regression framework.

4.2. Did the tick size war affect market quality?

Having documented that the tick size war had significant effects on both pre-trade and post-trade market shares, we next consider its effects on market quality, both at individual trading venues and in the aggregate.

Our empirical strategy utilizes a difference-in-differences design, which is motivated by two considerations. The first is that market quality on average in the period 1995–2008 is significantly worse in June than May—even absent a tick size war. Internet Appendix Table A.3 uses trading data from the "placebo" period 1995–2008 and shows that a simple before-after comparison of June 2009 and May 2009 is likely to conflate the effects of the tick size war with this seasonal (summer) effect. The second consideration is that the observed seasonal effect is almost identical for the most-traded stocks and for the less-traded stocks, making them excellent candidates to act as control stocks.

Combining these two considerations, our difference-in-differences model estimates the change in market quality from May to June 2009 for the most-traded stocks in the Scandinavian markets—that is, those traded outside the primary exchanges, and thus affected by the June, 2009 MTF tick reductions—and subtracts the trend for a control group of less-traded stocks not affected by the tick size reductions. Our baseline control sample contains all primary exchange stocks that were not traded on the MTFs throughout 2009, and which were as such not affected by the reductions in tick sizes on the MTFs. We also show the robustness of our results to using only the more liquid

Fig. 2: Effects of tick size war on market shares

This figure plots two measures of primary exchange market shares throughout the tick size war. In solid black, as our pre-trade market share measure, we plot the fraction of the trading day that the primary exchange quotes at the market-wide best bid and ask prices. In dashed red, as our post-trade market share measure, we plot the fraction of market-wide trading volume that is executed on the primary exchange. The plot only considers stocks in the blue-chip indexes at each of the primary exchanges. The gray vertical lines indicate the main events of the tick size war. The leftmost vertical line indicates June 1, 2009, the start of the tick size war. The rightmost vertical line indicates the tick size harmonization date, which varies across markets. For Norway and Sweden there is a middle line, which indicates the date when the primary exchange retaliated by lowering its tick sizes.

Panel A: Denmark (Copenhagen)

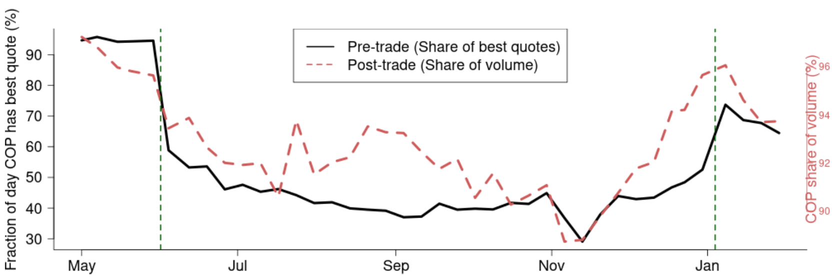

Panel B: Norway (Oslo)

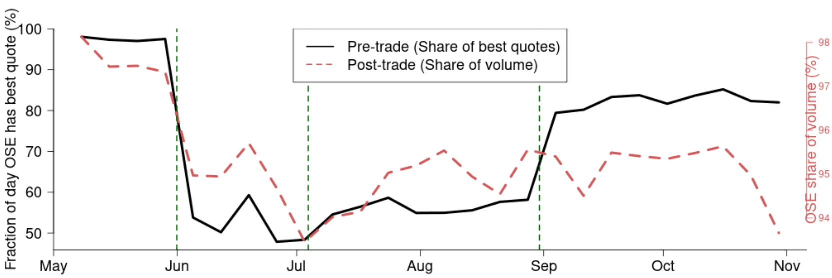

Panel C: Sweden (Stockholm)

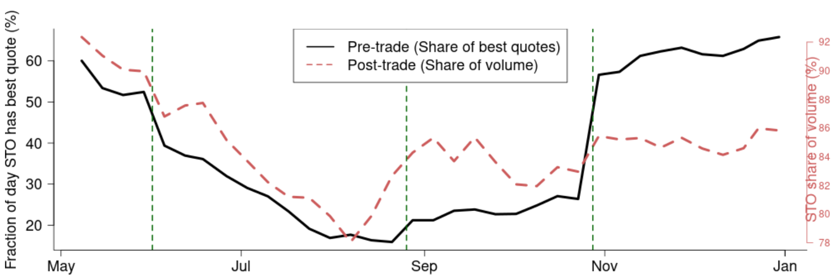

subset of these control stocks. We estimate the following regression:

$$y_{it} = \alpha_i + \alpha_t \times \alpha_c + \tau D_{it} + \omega_{it}, \tag{1}$$

where y_{it} is the outcome of interest, for example the effective spread, for stock i on date t; $D_{it} = Post_{it} \times Treated_{it}$ and hence equals one for stock i in the treated group ($Treated_{it} = 1$) after the start of the tick size war ($Post_{it} = 1$), and zero otherwise; α_i are stock-venue fixed effects; [14] and α_t are date-level fixed effects. The inclusion of α_i and α_t controls for fixed differences in y_{it} between treated and control stocks, and ensures that the before-after effect of the tick size war for treated stocks is measured net of the time trend for control group stocks. To allow for the possibility that stocks in different markets follow different time trends, we interact the date fixed effect α_t with α_c, a market-level fixed effect. Thus, our difference-in-differences model compares treated stocks in the Danish (or Norwegian or Swedish) market exclusively to control stocks from the same market. We estimate τ using a short sample period (May 1 – June 30) to exclude the primary exchanges' tick size retaliations. Equation (1) is estimated separately for stocks traded on the primary exchanges and the MTFs.[15]

The difference-in-differences estimates are presented in Table 3. We transform our spread, depth, and volume measures with the natural logarithm; for log-transformed variables, the difference-in-differences estimate, τ, multiplied by 100 can approximately be interpreted as the percentage effect of the tick size war. For the primary exchanges, we find economically and statistically significant reductions of about 8% for both the quoted and effective spreads as a result of the tick size war. The latter effect is primarily driven by a reduction in the realized spread. We also find a 12% increase in primary exchange volume. For the alternative venues, we observe more pronounced improvements in measured liquidity. These include significant reductions in quoted spread of almost 33%, which highlight the large reductions in spreads that are made possible by tick

[14] Throughout the paper, the stock identifier i corresponds to the unique stock-venue combination. For example, Statoil traded on the OSE (STL.OL) is treated as a separate unit i from Statoil traded on Chi-X (STL.CHI). This separation allows us to explore how market quality in the same stock was affected by the tick size war separately based on the venue of trading. When considering aggregate measures of market quality, such as the NBBO bid-ask spread, Statoil is represented by the same NBBO for both primary and MTFs venues. In such cases, the stock-venue fixed effect collapses to a simple stock fixed effect.

[15] Because our difference-in-differences model includes stock-venue fixed effects, the estimation sample effectively retains only the stocks that have trading activity both before and after the start of the tick size war. To avoid contamination from index inclusion and exclusion effects, we exclude from the estimation sample stocks that enter or exit the Copenhagen, Oslo, or Stockholm blue-chip indexes during 2009.

size changes. These reduced quoted spreads generate measurable benefits for those demanding liquidity, reducing the effective spreads paid by 27%. Again, the majority of these improvements come at the cost of the liquidity suppliers, with realized spreads declining significantly. The increases in volume on the alternative exchanges are also more pronounced, increasing by over 66%.

To empirically assess the effects of the tick size war on aggregate market quality, we re-estimate Equation (1) using measures of market quality based on the NBBO, together with total trading volume across all venues, as the outcome variables. The results are presented in the “Aggregate” columns of Table 3. The results show a significant reduction in NBBO spreads, and significant increases in both total quoted depth and total traded volume, confirming that the tick size war had beneficial effects for overall market quality.

Table 3. Effects of the tick size war on market quality

The table reports estimates of the difference-in-differences effect τ from $y_{it} = \alpha_i + \alpha_t \times \alpha_c + \tau D_{it} + \omega_{it}$. The sample period is May 2009 through June 2009. The effect τ is estimated separately for primary exchanges (OSE, STO, COP), MTFs (Chi-X, Turquise, BATS), and the NBBO. In our main specification, labeled "Full Sample," the treated sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009, and the control sample comprises all stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. In the "Traded $\geq$ 70%" and "Traded $\geq$ 90%" robustness specifications, we restrict the control sample to stocks that are traded on at least 70% (90%) of the trading days in 2009. In the "Propensity Matched" robustness specification, we use a one-to-one nearest neighbor matching procedure to restrict the control sample to the 89 control stocks that are the closest matches to our 89 treated sample stocks based on average bid-ask spreads, depth, and trading volume during the period from January to May 2009. All dependent variables apart from realized volatility are transformed with the natural logarithm. All regressions control for the natural logarithm of the stock price. For ease of exposition, we have included N, the number of observations, only for regressions using trading volume as the dependent variable. The number of observations varies slightly across dependent variables. t-statistics are presented in parentheses. Standard errors are clustered at the stock-venue level. Significance levels are indicated as: ** $p <$ 10%, ** $p <$ 5%, *** $p <$ 1%.*

	Full Sample			Robustness Tests: Traded $\geq$ 70%			Robustness Tests: Traded $\geq$ 90%			Robustness Tests: Propensity Matched		
	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate
Quoted spread	-0.08^{***}	-0.33^{***}	-0.20^{***}	-0.08^{***}	-0.34^{***}	-0.20^{***}	-0.09^{***}	-0.33^{***}	-0.21^{***}	-0.09^{***}	-0.26^{***}	-0.20^{***}
	(−4.40)	(−12.30)	(−8.39)	(−4.50)	(−12.35)	(−8.55)	(−4.44)	(−12.20)	(−8.58)	(−3.87)	(−9.02)	(−7.56)
Effective spread	-0.08^{***}	-0.27^{***}	-0.05^{***}	-0.08^{***}	-0.27^{***}	-0.05^{**}	-0.07^{***}	-0.26^{***}	-0.04^{**}	-0.08^{***}	-0.21^{***}	-0.06^{**}
	(−3.99)	(−10.73)	(−2.62)	(−3.78)	(−10.58)	(−2.41)	(−3.42)	(−10.27)	(−2.07)	(−3.32)	(−7.25)	(−2.26)
Realized spread	-0.15^{***}	-0.31^{***}	−0.04	-0.15^{***}	-0.31^{***}	−0.03	-0.15^{***}	-0.31^{***}	−0.03	-0.13^{***}	-0.22^{***}	−0.03
	(−3.93)	(−7.29)	(−1.22)	(−3.78)	(−7.18)	(−1.04)	(−3.71)	(−7.15)	(−0.97)	(−3.00)	(−4.62)	(−0.67)
Price impact	−0.04	-0.24^{***}	-0.18^{***}	−0.04	-0.23^{***}	-0.17^{***}	−0.03	-0.22^{***}	-0.16^{***}	−0.04	-0.19^{***}	-0.17^{***}
	(−1.40)	(−5.66)	(−4.59)	(−1.24)	(−5.58)	(−4.45)	(−0.88)	(−5.30)	(−4.09)	(−0.97)	(−3.86)	(−3.52)
Depth	0.00	−0.00	0.11^{***}	0.00	−0.01	0.11^{***}	0.01	−0.01	0.11^{***}	0.01	0.03	0.11^{***}
	(0.10)	(−0.16)	(3.58)	(0.14)	(−0.31)	(3.60)	(0.22)	(−0.42)	(3.61)	(0.18)	(1.09)	(3.07)
Volatility	0.01	0.03	−0.01	0.02	0.04	−0.00	0.05	0.06^{**}	0.03	0.02	0.02	−0.00
	(0.10)	(1.21)	(−0.05)	(0.14)	(1.31)	(−0.00)	(0.44)	(2.14)	(0.29)	(0.12)	(0.27)	(−0.01)
Volume	0.12^{***}	0.66^{***}	0.16^{***}	0.13^{***}	0.66^{***}	0.17^{***}	0.16^{***}	0.67^{***}	0.19^{***}	0.15^{**}	0.55^{***}	0.19^{***}
	(2.65)	(13.68)	(3.40)	(2.77)	(13.56)	(3.49)	(3.08)	(13.38)	(3.78)	(2.23)	(9.12)	(2.73)
# treated units	89			89			89			89		
# control units	577			451			317			89		
N	23,344			20,328			15,820			6,987		

The validity of our difference-in-differences estimates in Table 3 does not require that treated and control sample stocks are similar in their levels of liquidity, only that they follow similar trends. Our primary specification, reported in the "Full Sample" panel of Table 3, uses 89 treated stocks, as well as 577 control stocks traded across the Scandinavian exchanges that did not experience any tick size competition.[16] Internet Appendix Table A.4 shows that although these treated and control stocks significantly differ in their levels of liquidity, their trends in liquidity are statistically indistinguishable in the pre-war period January–May, 2009, supporting the common trends assumption of the difference-in-differences design. Internet Appendix Figure A.7 provides graphical evidence of the same effect. In addition, Internet Appendix Table A.3 shows evidence of common trends in the 14 years 1995–2008 leading up to the tick size war.[17]

For transparency and robustness, however, we explore the sensitivity of our difference-in-differences estimates to using three alternative, more liquid samples of control stocks. The first two robustness tests progressively exclude more control stocks based on their number of days with trading activity throughout 2009. In the panel labeled "Traded $\geq$ 70%," we retain only the 451 control stocks that traded on at least 70% of the approximately 250 trading days; the panel "Traded $\geq$ 90%" extends this requirement to at least 90% of the trading days in 2009, resulting in a further reduction to 317 control stocks. Going even further, in the panel "Propensity matched," we keep only the 89 control stocks that most closely match the liquidity profile of our 89 treated stocks. To achieve this, we undertake a one-to-one nearest neighbor matching based on spreads, depth and volume in the period January–May, 2009. Across all three robustness tests, estimates are virtually unchanged compared to the primary specification, all suggesting that the tick size war benefited market quality.

Given that our main results in Table 3 are virtually unaffected by the choice of control sample, the following sections will continue to focus on the "Full Sample" specification in Table 3, as that specification provides the most statistical power. The added statistical power is particularly useful in analyses where we partition the data to examine effect heterogeneity.

[16] The tick size war was shorter for 4 stocks traded on the Stockholm Stock Exchange (ALFA, LUPE, TEL2a, and TEL2b), with their primary listing tick sizes reduced on June 8, 2009. Our results are not sensitive to the inclusion or exclusion of these securities. Additionally, Nordea and Nokia, along with other Stockholm listings, traded at different tick sizes at the primary exchange and the MTFs even before the tick size war started on June 1, 2009.

[17] See Internet Appendix Section C for a comprehensive discussion of the difference-in-differences model.

4.3. Do binding tick sizes matter?

We have seen that the tick size war during June 2009 caused a reduction in the aggregate trading costs and an increase in the aggregate trading volume. To better understand the mechanisms driving these results, we assess whether the effects vary with the extent to which individual stocks' bid-ask spreads were constrained by the tick size in the pre-war period. There are two main reasons why we expect the effects to vary by the extent of spread constraints. In stocks for which the tick size is more frequently binding, permitting a finer pricing grid provides an opportunity for liquidity suppliers to reduce quoted bid-ask spreads closer to their "true" cost of liquidity provision. Such improvements are likely to significantly benefit market quality, increasing the available depth at (or better than) the primary exchange's best quotes, and reducing trading costs for all market participants.

In contrast, for securities where the bid-ask spread is unconstrained by the tick size, market participants have the opportunity to provide price improvement without requiring a finer price grid. Providing a finer price grid in these unconstrained trading environments allows liquidity providers to "undercut" the primary exchange quotes by a single tick, similar to providing fractional price improvement in the dark or "paying to post" on an inverted maker-taker venue.[18] Chung et al. (2020) find that a market-wide tick size *increase* results in higher costs for small orders, but generates benefits for larger orders. Providing a finer pricing grid on only some trading venues *may* improve measures of market quality such as the quoted and effective bid-ask spreads for small, marginal orders, but the predicted impacts for larger orders, order book depth, and market-wide market quality are less clear.

To shed light on this mechanism, we need to gauge the extent to which trading in a given stock is constrained by its current minimum tick size. We use two separate measures of spread constraints. The first is based on the much-used "leeway" metric.[19] In particular, for stock i on date t, we start by calculating:

$$\text{TicksPerSpread}_{\text{it}} = \frac{\text{QuotedSpread}_{it}}{\text{TickSize}_{it}},$$

As $\text{TicksPerSpread}_{\text{it}}$ decreases, the quoted spread comes closer to being constrained by the tick size. In the limit, where $\text{TicksPerSpread}_{\text{it}} = 1$, the tick size entirely constrains the quoted bid-ask spread, preventing it from decreasing. Then, we base

[18] Such undercutting behavior, experienced for example in the trading environment for cryptocurrencies, where tick sizes are extremely small, has been shown to increase measures of transactions costs, discouraging limit order provision (Dhyrberg et al., 2019).

[19] See, for example, Autorité des marchée financiers (AMF) (2018).

Table 4. Effects on market quality, based on binary constraint measure

The table reports estimates of the difference-in-differences effect τ from Equation (1). The sample period is May 2009 through June 2009. The effect τ is estimated separately for primary exchanges (OSE, STO, and COP), MTFs (Chi-X, Turquoise, and BATS), and the NBBO. The baseline treated sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The baseline control sample comprises all stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. We partition the treated sample into *Constrained* and *Unconstrained* stocks, as defined by the binary constraint measure described in Section 4.3, and estimate the difference-in-differences specification separately for each of the treated samples, holding fixed the definition of the control sample. All dependent variables apart from realized volatility are transformed with the natural logarithm. All regressions control for the natural logarithm of the stock price. For ease of exposition, we have included N, the number of observations, only for regressions using trading volume as the dependent variable. The number of observations varies slightly across dependent variables. t-statistics are presented in parentheses. Standard errors are clustered at the stock-venue level. Significance levels are indicated as: ** $p < 10\%$, ** $p < 5\%$, and *** $p < 1\%$.*

	Constrained			**Unconstrained**		
	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate
Quoted spread	−0.10***	−0.34***	−0.25***	−0.04	−0.29***	−0.08***
	(−4.74)	(−11.21)	(−8.85)	(−1.61)	(−5.01)	(−3.05)
Effective spread	−0.10***	−0.28***	−0.08***	−0.05*	−0.23***	−0.01
	(−4.02)	(−10.10)	(−3.44)	(−1.84)	(−3.92)	(−0.23)
Realized spread	−0.18***	−0.32***	−0.06	−0.10*	−0.28**	−0.03
	(−3.85)	(−6.92)	(−1.47)	(−1.87)	(−2.54)	(−0.53)
Price impact	−0.05*	−0.24***	−0.23***	−0.02	−0.28***	−0.06
	(−1.75)	(−5.04)	(−5.42)	(−0.37)	(−3.40)	(−1.15)
Depth	0.02	−0.02	0.13***	−0.02	0.09**	0.09
	(0.52)	(−0.74)	(3.93)	(−0.43)	(2.21)	(1.58)
Volatility	0.09	0.03	0.06	−0.14	0.04	−0.14
	(0.64)	(1.12)	(0.45)	(−0.79)	(0.91)	(−0.78)
Volume	0.17***	0.64***	0.22***	0.04	0.75***	0.05
	(3.96)	(12.40)	(5.06)	(0.41)	(7.22)	(0.61)
# treated units	55			34		
# control units	577			577		
N	22,018			21,172		

a binary categorization on the average of $\text{TicksPerSpread}_{\text{it}}$ on the primary exchange one month prior to the tick size war. Stocks with $\text{Mean}(\text{TicksPerSpread}_{\text{it}}) \geq 2$ are defined as *Unconstrained*, while stocks with $\text{Mean}(\text{TicksPerSpread}_{\text{it}}) < 2$ are defined as *Constrained*.

In Table 4 we re-estimate the difference-in-differences effect τ from Equation (1) separately for treated stocks classified as *Constrained* and *Unconstrained*. We find that most of the observed market quality effects are driven by the *Constrained* stocks, with statistically and economically significant reductions of about 10% in both quoted and effective bid-ask spreads on primary exchanges, accompanied by even larger spread reductions (around 30%) on alternative venues. For *Unconstrained* stocks, by contrast, spreads only decline on the alternative venues. Similarly, we find increased volume on both primary (+17%) and alternative (+64%) venues for *Constrained* stocks, whereas for *Unconstrained* stocks we only find a significant increase in trading volume on the alternative venues (+75%).

In columns labeled "Aggregate," Table 4 also examines the role of binding tick sizes on our measures of aggregate market quality. Consistent with the results for the individual venues, we find stronger effects on aggregate market quality for Constrained than Unconstrained stocks. For the quoted NBBO, we find reductions of 25% and 8% for *Constrained* and *Unconstrained* stocks, respectively. We also observe a statistically significant reduction in the NBBO effective spread (−8%) for Constrained stocks and a statistically and economically insignificant effect for Unconstrained stocks (−1%). Similarly, the observed increase in aggregate trading volume is almost entirely driven by a 22% increase in trading volume in Constrained stocks, compared with the statistically insignificant 5% increase for Unconstrained stocks. Overall, the results using the binary constraint measure suggest that the vast majority of the benefits conferred by the tick size war are attributable to those securities whose bid-ask spreads were previously constrained by the minimum tick size.

Our second measure of tick size constraints is based on a procedure developed by Kwan et al. (2015) to estimate the "true" bid-ask spread. Although the binary constraint measure used in Table 4 may separate Unconstrained stocks from Constrained stocks, *within* the set of Constrained stocks, the binary measure does not allow for a separation of stocks that are *heavily* constrained from others for which one tick is about right. This occurs owing to the lower bound on quoted spreads of one tick. As such, stocks that are heavily Constrained may be more likely to benefit from tick size competition. Failing to adequately identify these stocks will make it difficult to fully quantify the observed benefits. Following the methodology developed by Kwan et al. (2015), we construct our second measure of spread constraints in four steps. In the first step, we run the following regression for stock i on date t:

$$\ln(\text{QuotedSpread}_{it}) = \beta_0 + \beta_1 \ln(\text{Mcap}_{it}) + \beta_2 \ln(\text{Price}_{it}) + \beta_3 \ln(\text{Ntrades}_{it}) + \beta_4 \ln(\text{TradeSize}_{it}) + \beta_5 \text{Volatility}_{it} \quad (2)$$

where Mcap_{it} is the stock's market capitalization, Price_{it} is the end-of-day stock price, Ntrades_{it} is the number of trades, TradeSize_{it} is the average trade size, and Volatility_{it} is the realized volatility. We estimate Equation (2) on a sample of primary exchange stock-days prior to the tick size war where the spread was unconstrained by the tick size (as measured using $\text{TicksPerSpread}_{\text{it}}$). Then, in the second step, we use the estimated

βs from Equation (2), along with the observed daily levels of stock covariates, to predict $\ln(\widehat{\text{QuotedSpread}}_{it})$ for all the stock-days in our sample. In the third step, we calculate a continuous measure of spread constraints by comparing the observed spread to that implied by our model,

$$Constrained_{it}^{Cont.} = \ln(\text{QuotedSpread}_{it}) - \ln(\widehat{\text{QuotedSpread}}_{it}). \tag{3}$$

In Equation (3), positive (negative) values of $Constrained_{it}^{Cont.}$ suggest that trading in stock i on date t is spread constrained (unconstrained), as the actual QuotedSpread_{it} exceeds (is lower than) its predicted value, $\widehat{\text{QuotedSpread}}_{it}$. Hence, stocks with a large value of $\text{TicksPerSpread}_{\text{it}}$ may still be defined as *Constrained*. For example, if a stock-day sees $\widehat{\text{QuotedSpread}}_{it} = 1$, $\text{QuotedSpread}_{it} = 2$, and $\text{TickSize}_{it} = 1$, this stock-day will still be defined as Constrained by Equation (3), even though $\text{TicksPerSpread}_{\text{it}} = 2$. In the final step, we create three constraint categories based on $Constrained_{it}^{Cont.}$:

- Constrained: $Constrained_{it}^{Cont.} > 0.15$,
- At-margin: $-0.15 \leq Constrained_{it}^{Cont.} \leq 0.15$,
- Unconstrained: $Constrained_{it}^{Cont.} < -0.15$,

where the cutoffs are chosen to retain similar-sized stock samples within each category.[20]

[20] Section D of the Internet Appendix provides the distribution of $Constrained_{it}^{Cont.}$. In their empirical analysis, Kwan et al. (2015) linearly interact the continuous variable $Constrained_{it}^{Cont.}$ with their treatment of interest. Such a procedure assumes a linear functional form on the relationship between the treatment of interest and $Constrained_{it}^{Cont.}$. To circumvent the need for such functional form assumptions, we estimate treatment effects separately within each of the above bins, which allows for treatment effects that vary more flexibly across categories of $Constrained_{it}^{Cont.}$.

Table 5. Effects on market quality, based on continuous constraint measure

The table reports estimates of the difference-in-differences effect τ from Equation (1). The sample period is May 2009 through June 2009. The effect τ is estimated separately for primary exchanges (OSE, STO, and COP), MTFs (Chi-X, Turquoise, and BATS), and the NBBO. The baseline treated sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The baseline control sample comprises all stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. We partition the treated sample into *Constrained*, *At-margin* and *Unconstrained* stocks, as defined by the Kwan et al. (2015) procedure described in Section 4.3, and estimate the difference-in-differences specification separately for each of the treated samples, holding fixed the definition of the control sample. All dependent variables apart from realized volatility are transformed with the natural logarithm. All regressions control for the natural logarithm of the stock price. For ease of exposition, we have included N, the number of observations, only for regressions using trading volume as the dependent variable. The number of observations varies slightly across dependent variables. t-statistics are presented in parentheses. Standard errors are clustered at the stock-venue level. Significance levels are indicated as: * $p < 10\%$, ** $p < 5\%$, and *** $p < 1\%$.

	Constrained			**At-margin**			**Unconstrained**		
	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate	Primary	MTFs	Aggregate
Quoted spread	−0.10***	−0.43***	−0.26***	−0.06***	−0.26***	−0.15***	−0.05*	−0.30***	−0.05*
	(−3.74)	(−10.29)	(−7.56)	(−2.69)	(−6.90)	(−4.96)	(−1.67)	(−3.25)	(−1.92)
Effective spread	−0.12***	−0.36***	−0.10***	−0.04*	−0.19***	−0.00	−0.06**	−0.27***	−0.04
	(−3.90)	(−9.77)	(−4.01)	(−1.81)	(−5.53)	(−0.15)	(−2.33)	(−3.12)	(−1.24)
Realized spread	−0.17***	−0.42***	−0.07	−0.12**	−0.24***	0.01	−0.20***	−0.09	−0.12**
	(−3.10)	(−6.83)	(−1.51)	(−2.04)	(−3.86)	(0.22)	(−3.13)	(−0.83)	(−2.17)
Price impact	−0.09**	−0.33***	−0.25***	−0.01	−0.16***	−0.12**	0.03	−0.24	−0.00
	(−2.36)	(−5.41)	(−5.07)	(−0.21)	(−2.79)	(−2.47)	(0.58)	(−1.55)	(−0.02)
Depth	−0.02	−0.10***	0.05	0.04	0.09***	0.17***	−0.03	0.08	0.15*
	(−0.33)	(−3.00)	(1.07)	(1.32)	(3.20)	(5.17)	(−0.36)	(1.32)	(1.73)
Volatility	0.35**	0.03	0.31*	−0.23**	0.03	−0.22**	−0.59	0.06	−0.58
	(2.13)	(1.09)	(1.93)	(−2.01)	(0.88)	(−1.99)	(−1.27)	(1.10)	(−1.26)
Volume	0.13**	0.63***	0.16***	0.13**	0.70***	0.16**	0.06	0.70***	0.09
	(2.39)	(10.26)	(3.02)	(2.01)	(9.83)	(2.57)	(0.39)	(5.23)	(0.59)
# treated units	42			36			11		
# control units	577			577			577		
N	21,492			21,268			20,276		

Table 5 presents estimates from our difference-in-differences estimations grouping the treated stocks into the three categories defined by the Kwan et al. (2015) measure, *Constrained*, *At-margin*, and *Unconstrained*. The results corroborate our previous findings of stronger market quality improvements as the spread constraint becomes more binding. Across nearly all venue-specific and aggregate measures of market quality, the effect sizes τ are monotonically decreasing in the extent of spread constraints. For example, considering the NBBO quoted spread, we find reductions of 26%, 15%, and 5% for *Constrained*, *At-margin*, and *Unconstrained* stocks, respectively. Similarly, we find increases in trading volume for *Constrained* and *At-margin* stocks, but no volume effects for *Unconstrained* stocks.

4.4. What happens to market-wide informational efficiency?

Our results in Sections 4.2 and 4.3 show that the tick size competition that took place during June 2009 significantly improved measures of bid-ask spreads, depth, and trading volume, particularly for the most constrained stocks. To provide a more complete picture of the effect of tick size competition on overall market quality, we now turn to the impact of the tick size war on the informational efficiency of the (overall) market place.

Our analysis of informational efficiency builds on a sequence of papers by Chordia et al. (2005, 2008, 2011). These papers link informational efficiency to the horizon over which it is estimated. Even if stock prices are efficient—that is, follow a random walk—over a daily horizon, intraday stock returns may exhibit systematic predictability. Chordia et al. (2005) argue that intraday deviations from random walks can be used to measure the extent of informational efficiency: the closer the midpoint returns are to a random walk, the more efficient the price process. Chordia et al. (2005) analyze the US stock market across its two major tick size changes, from eights to sixteenths to decimals, and provide compelling evidence that price efficiency improves when tick sizes narrow.

In our setting, we want to investigate whether adding the option to quote at narrower tick sizes at the satellite exchanges improves market-wide informational efficiency. We therefore first aggregate across markets and calculate the midquote between the best prices: our NBBO estimates. Based on these NBBO midquotes, we then calculate two standard measures of informational efficiency: the return autocorrelation, and the variance ratio of Lo and MacKinlay (1988).

The two efficiency measures for stock i are calculated from returns r_i as follows:

$$\text{Autocorrelation:} \quad AC_{i,k} = \left| Corr\left(r_{i,\tau}, r_{i,\tau+k}\right)\right|. \tag{4}$$

$$\text{Variance Ratio:} \quad VR_{i,k,l} = \left| \frac{\sigma^2_{i,k}}{l\sigma^2_{i,kl}} - 1 \right|. \tag{5}$$

Here, k is a time interval, l is an integer, and τ is an observed time. If, for example, k is one minute, $AC_{i,k}$ is the autocorrelation of one minute returns. For the variance ratio, $\sigma^2_{i,k}$ is the variance of k- or lk-period returns. For robustness, we consider several frequencies for both of these measures. For $AC_{i,k}$, we consider ten-second and one-minute frequencies. For $VR_{i,k,l}$, we consider variance ratios of 10 seconds versus 50 seconds and 1 minute versus 5-minutes. These choices follow Foley and Putniņš (2016). We calculate $AC_{i,k}$ and $VR_{i,k,l}$ for each stock on a daily basis. To deal with potential bias in the construction of our measures in the presence of infrequent trading, we exclude stock-days with less than 10 transactions. For both $AC_{i,k}$ and $VR_{i,k,l}$ the economic interpretation is that the closer the measure is to zero, the more informationally efficient is the market.

Table 6 presents the difference-in-differences estimates of τ from Equation (1) using $AC_{i,k}$ and $VR_{i,k,l}$ as the outcomes. Consistent with our results in Tables 3–5, the tick size war has an insignificant effect on the informational efficiency of the overall market. However, when securities are split by their level of constraint, a more nuanced picture emerges. Although we find no significant effect on efficiency for the unconstrained securities, constrained securities exhibit significant improvements in their levels of informational efficiency. Both $AC_{i,k}$ and $VR_{i,k,l}$ decrease significantly, irrespective of the measurement period. These results are consistent with market makers using the finer pricing grid to more accurately price securities, particularly when trading was previously constrained by the tick size.

4.5. Did the primary exchanges recapture market share?

We have so far focused on the *breakout phase* of the tick size war. We now turn to the *retaliation phase*, in which Oslo and Stockholm retaliated to the MTFs' tick size reductions by reducing the tick size for some of their most heavily traded stocks.[21] We are particularly interested in whether the primary exchanges were able to *recapture* market share from the MTFs by reducing their own tick size. For each retaliation event, we identify for each treated stock the most comparable non-treated blue-chip index stock from either Oslo, Copenhagen, or Stockholm using propensity score matching based on spreads, depth, and volume.[22] Then, we estimate a difference-in-differences model

[21] As shown in Table 1, on July 6, 2009, Oslo lowered its tick size for 25 heavily-traded stocks, and on August 26, 2009, Stockholm lowered its tick size for 10 of its most heavily traded stocks.

[22] We restrict the control sample to index stocks because the primary exchange market share is close to unity and varies little for stocks outside the blue-chip indexes. We note that a similar strategy of comparing treated index stocks to non-treated index stocks would not

Table 6. Effects on NBBO pricing efficiency

The table reports estimates of the difference-in-differences effect τ from Equation (1). The sample period is May 2009 through June 2009. The effect τ is estimated only for the NBBO. As our dependent variables, we consider two measures of NBBO informational efficiency: 10 second and 1 minute autocorrelations, and two variance ratios, 10 second/50 second and 1 minute/5 minute. The measures are defined in Equations (5)–4), all calculated from (NBBO) midquote returns. When calculating the variance ratio and autocorrelation of NBBO returns, we restrict attention to trading days with at least 10 transactions. In the "Full Sample" column, the treated sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009, and the control sample comprises all stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. In the "Constrained" and "Unconstrained" columns, we partition the treated sample into *Constrained* and *Unconstrained* stocks, as defined by the binary constraint measure described in Section 4.3, and estimate the difference-in-differences specification separately for each of the treated samples, holding fixed the definition of the control sample. All regressions control for the natural logarithm of the stock price. For ease of exposition, we have included N, the number of observations, only for regressions using the 50-second variance ratio as the dependent variable. The number of observations varies slightly across dependent variables. t-statistics are presented in parentheses. Standard errors are clustered at the stock-venue level. Significance levels are indicated as: $^{*}\ p < 10\%$, $^{**}\ p < 5\%$, and $^{***}\ p < 1\%$.

	Full sample	Constrained	Unconstrained
Efficiency measure			
$AC_{1\text{ min}}$	−0.005	−0.008**	0.002
	(−1.38)	(−1.98)	(0.26)
$AC_{10\text{ s}}$	−0.003	−0.005*	−0.000
	(−1.34)	(−1.89)	(−0.01)
$VR(1\text{ min}, 5\text{ min})$	−0.009	−0.015*	0.003
	(−1.19)	(−1.86)	(0.24)
$VR(10\text{ s}, 50\text{ s})$	−0.006	−0.013**	0.007
	(−0.97)	(−2.04)	(0.71)
N	13,578	12,253	11,406

comparing market shares for treated and control stocks before and after the retaliation event.

The results are presented in Table 7. Despite significant reductions in their own tick size, the primary exchanges are unable to recapture the market shares of executed

be feasible in the difference-in-differences analysis of the *breakout phase*. This is because all the blue-chip index stocks in Stockholm, Oslo, and Copenhagen, as well as some non-index stocks, were traded on the MTFs and thus affected by the MTF tick size reductions, leaving no unaffected index stocks as eligible controls. Hence, the analyses in Tables 3–6 compared stocks with (treated) and without (control) MTF trading, before and after the *breakout phase*.

volume from the MTFs. This evidence supports the existence of a fragmented competitive equilibrium , consistent with the findings of Chen and Duffie (2021), and the observed failure of primary exchanges to recapture lost volume is consistent with the existence of liquidity externalities in exchange competition (e.g., Pagano (1989)). The results in Table 7 are also consistent with existing evidence of significant fixed costs to venue connectivity, as documented by Aitken et al. (2017). In particular, once a fixed connectivity cost has been paid, the marginal benefits of connectivity likely outweigh the marginal costs, resulting in continued broker connection.[23] For market shares of *quoted* volume, however, Table 7 shows evidence that the primary exchanges were able to recapture market shares in constrained securities. This is consistent with quote competition on the primary exchange from brokers who had refrained from connecting to the MTFs.

Table 7 also reports the effects of the Stockholm and Oslo retaliation events on measures of market quality at the primary exchanges. The results for market quality are consistent with the majority of existing evidence on the effects of tick size reductions (e.g., Bessembinder (2003)). In particular, we find economically and statistically significant reductions in trading costs and depth for constrained stocks, but no effects for unconstrained stocks.

5. Order Submission Strategies

Our main results in Section 4 suggest that the tick size war of June 2009 disproportionately improved the market quality of stocks where the spread was constrained by the pre-existing tick size regime. Such improvements suggest that a binding tick size can increase transactions costs for active, liquid securities. Allowing for a finer pricing grid on a single trading venue was, in our setting, sufficient to encourage competition for order flow across venues, reducing overall trading costs, as well as increasing aggregate volume and quoted depth. This stands in stark contrast to the sample of unconstrained stocks, for which we observe few statistically significant effects, indicating that the finer pricing grid was not necessary.

To better understand the mechanisms underlying these results, in this section, we analyze the order submission strategies of liquidity providers in the presence of both a small- and large-tick exchange, documenting the level of price improvement provided, and gauge whether undercutting is more likely than true price competition.

[23] Similar findings are presented for the Canadian market by Foley et al. (2019): the removal of the requirement to connect to existing venues did not result in a reduction of broker connectivity.

Table 7. Did the primary exchanges recapture market share?

The table reports the average difference-in-differences effect on primary exchange market shares and market quality of the OSE and STO tick size retaliation events on July 6, 2009, and August 26, 2009, respectively. The estimates are obtained in four steps. First, for each retaliation event, we compile a list of treated stocks affected by the retaliation (see Table 1). Second, we use nearest neighbor matching to find a single control stock for each treated stock. We match based on average spreads, depth, and volume in the 10 trading days leading up to the retaliation event, and we restrict the matching to blue-chip index stocks at the OSE, STO, and COP. For the STO retaliation event, we restrict the matching to index stocks at the STO and COP because the OSE harmonized tick sizes on August 31, 2009, shortly after the STO event. Third, for each treated and matched control, we keep 10 trading days of data from before and after each retaliation event. Finally, we pool the data from the two retaliation events to estimate the following difference-in-differences specification: $y_{it} = \alpha_{ip} + \alpha_{tp} + \beta D_{it} + \omega_{it}$, where α_{ip} and α_{tp} are stock and time fixed effects, respectively. The stock and time fixed effects are retaliation event-specific, hence the p subscript, which ensures that the difference-in-difference effect β is estimated entirely based on variation *within* a given retaliation event (for example, treated stocks for the OSE retaliation event are only compared to matched control stocks for the OSE event). In the "Full Sample" column, the estimation sample comprises all treated and matched control stocks. In the "Constrained" and "Unconstrained" columns, the treated sample is restricted to *Constrained* and *Unconstrained* stocks, respectively, as defined by the binary constraint measure described in Section 4.3. All dependent variables apart from the two market share measures and realized volatility are transformed with the natural logarithm. All regressions control for the natural logarithm of the stock price. For ease of exposition, we have included N, the number of observations, only for regressions using trading volume as the dependent variable. The number of observations varies slightly across dependent variables. t-statistics are presented in parentheses. Standard errors are clustered at the stock-venue level. Significance levels are indicated as: * $p < 10\%$, ** $p < 5\%$, and *** $p < 1\%$.

	Full Sample	Constrained	Unconstrained
Market share of volume	0.50	0.62	−0.16
	(0.98)	(1.10)	(−0.24)
Market share of best quotes	8.03***	10.29***	2.82
	(3.29)	(3.89)	(0.80)
Quoted spread	−0.21***	−0.27***	−0.01
	(−4.08)	(−4.87)	(−0.13)
Effective spread	−0.34***	−0.41***	−0.10
	(−5.82)	(−6.68)	(−1.23)
Realized spread	−0.52***	−0.62***	−0.10
	(−4.00)	(−4.21)	(−0.59)
Price impact	−0.18**	−0.24***	−0.03
	(−2.54)	(−3.17)	(−0.22)
Depth	−0.65***	−0.71***	−0.38***
	(−6.74)	(−6.40)	(−3.88)
Volatility	0.56	0.60	0.07
	(1.08)	(1.14)	(0.11)
Volume	0.04	0.06	0.04
	(0.64)	(0.80)	(0.27)
N	1240	1080	780

5.1. Undercutting or price competition?

Fig. 3: Trading at large- and small-tick exchanges

The figure illustrates the differences in pricing grids at a large-tick primary exchange (top) and a small tick alternative venue (bottom). We expect traders seeking only to undercut prices at the primary exchange to improve the best quotes by only one small tick. True price competition would see more activity deeper within the alternative venue's (more granular) order book.

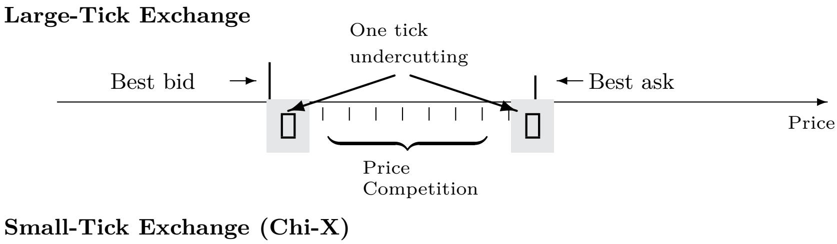

When competing venues allow trading with different pricing grids, liquidity providers have the opportunity to either undercut the large-tick exchange or aggressively compete with each other for order flow (see Figure 3 for an illustration). In our setting, the observed reductions to the NBBO spread could be driven by such undercutting behavior, where liquidity providers use the new finer pricing grid to undercut the primary exchange (as in Biais et al. 2010) by an "infinitesimal" amount. This behavior is predicted by among others the Foucault et al. (2005) and Bhattacharya and Saar (2020) models, where the choice of undercutting the best limit order is a trade-off between the benefits of being first in the order book queue and the cost of acquiring that position. Such costs are decreasing in the tick size, with the benefits of prime queue position likely accentuated when tick sizes are binding. Accordingly, we may observe single tick undercutting, where tick sizes were not particularly binding (as the cost of doing so was already relatively low), as well as multi-tick undercutting when costs are low and the value of being first is high (i.e., spreads are still comparatively wide).[24]

To empirically assess how liquidity providers exploit the access to a new finer pricing grid, we zoom in on the competitive dynamics between the most active alternative trading venue, Chi-X, and the primary exchanges. Figure 4 plots the distribution of the best quotes on Chi-X relative to the primary exchanges, measured in alternative-venue tick sizes, one week before and after the start of the tick size war. In the figure, "equal"

[24] Empirically, Battalio et al. (2016) demonstrate that undercutting behavior (in their setting, facilitated via inverted make-take venues) can have a negative "cream-skimming" effect on overall market quality. On the other hand, a finer pricing grid may encourage more active competition between market makers, which can cause various spread-based measures of transaction costs to narrow.

indicates that Chi-X and the primary exchange are quoting identical prices, whereas +1 (−1) indicates that the prices on Chi-X are one new tick better (worse) than the primary exchange.[25]

The results in Panels A–F of Figure 4 show a striking change in Chi-X traders' order submission strategies from just before to after the start of the tick size war, particularly for stocks listed in Copenhagen and Oslo. For these exchanges, just before the war, Chi-X quotes matched the best primary exchange prices about 50% of the time, and were more likely to be behind the best quotes than improving on them, suggestive of a less-active venue where relatively stale prices were quoted with minimal risk. Immediately after the start of the tick size war, however, Chi-X prices *improved* on the best prices about 60% of the time. The typical Chi-X price improvement was a single tick (35%-45% of the time), with multiple-tick price improvements consistent with "true" price competition occurring only 5%-10% of the time. These findings suggest that the opportunity to engage in single-tick undercutting was a key driver of the observed increase in Chi-X price improvement during the tick size war.

The results in Figure 4 for STO listed stocks tell a slightly different story. Unlike COP and OSE, Chi-X offered smaller tick sizes than STO even before the tick size war. For STO listings, the tick size war simply widened pre-existing tick size differences. Consequently, just before the tick size war, the Chi-X quotes for stocks in Stockholm were similar to the primary exchange quotes around 25% of the time, with single-tick price *improvements* 35% of the time (i.e., similar to the *post*-war distribution for Oslo and Copenhagen). After Chi-X further reduced their tick sizes for Stockholm stocks, Panel F of Figure 4 shows, Chi-X continued to offer one-tick price improvements 35% of the time but, in addition, began to offer *multiple*-tick price improvements over the primary exchange more than 20% of the time. These findings suggest that price competition on Chi-X was a key driver of the increase in price improvement for Stockholm stocks.

[25] The calculation underlying the statistics in Figure 4 is done separately for the bid and ask sides of the limit order book and then averaged. For the bid (ask) side, a price improvement means that the Chi-X price is above (below) the corresponding price on the primary exchange.

Fig. 4: Chi-X quoting behavior around the tick size reduction

The figure compares quoting behavior at Chi-X and the primary exchange before and after Chi-X's June 1 tick size reduction. For each COP, OSE, and STO listed stock with Chi-X trading, we calculate the fraction of the trading day that the Chi-X price is equal to the primary exchange price; improves on the primary exchange price by one, two, three, four, or five and more ticks; or is worse than the primary exchange price by one, two, three, four, or five and more ticks. For the bid (ask) side, a price improvement means that Chi-X quote prices that are above (below) the corresponding primary exchange quotes. The histograms present the average of these fractions across the sampled stock-days. We only consider stocks in the blue-chip indexes at COP, OSE, and STO.

One week pre-war

Panel A. Denmark

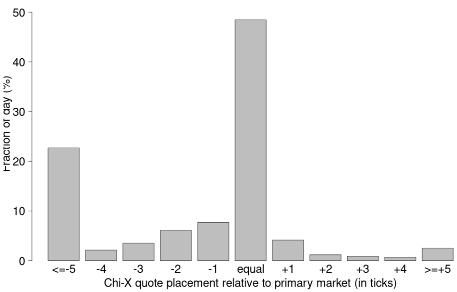

Panel B. Norway

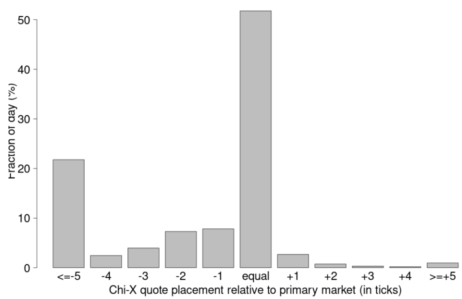

Panel C. Sweden

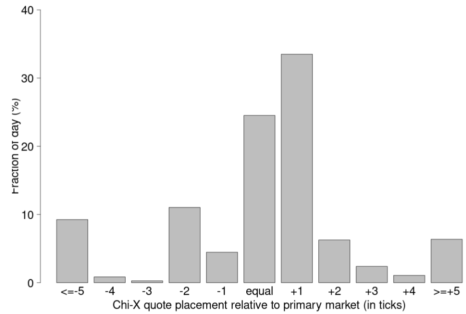

One week post-war

Panel D. Denmark

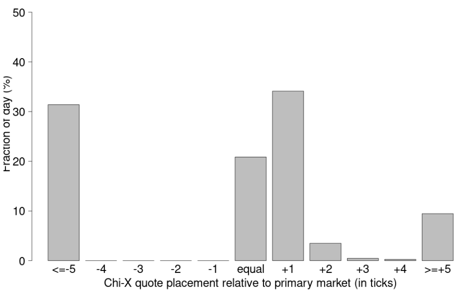

Panel E. Norway

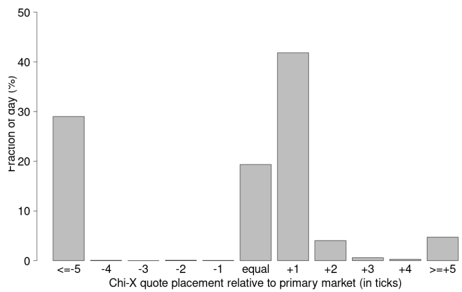

Panel F. Sweden

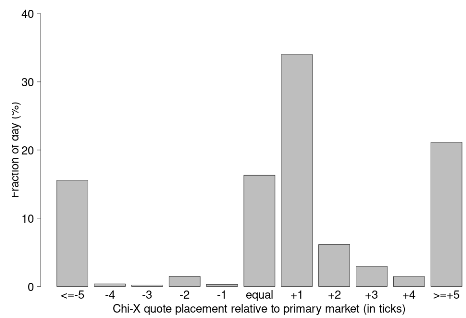

Overall, the results are consistent with the theoretical predictions of Bhattacharya and Saar (2020)—reducing the Chi-X tick size even further made the cost of undercutting even lower, resulting in more frequent undercutting, and hence price improvement. Our results are also consistent with empirical findings in both equity and cryptocurrency markets. O'Hara et al. (2018) show that market makers behave more cautiously when tick sizes are increased in the US tick size pilot—this implies that should the tick sizes narrow, market makers' limit orders would become more aggressive—which is exactly what we document. In the cryptocurrency arena, tick sizes orders of magnitude smaller than those observed in our study are commonplace, and Dhyrberg et al. (2019) show that these tiny tick sizes encourage similar competitive undercutting behavior.

5.2. War and peace: The role of competitive tick sizes

Figure 4 presents a single snapshot of Chi-X traders' order submission strategies just before the start of the tick size war and a single snapshot just after. To complement these snapshots, we present results for the full duration of the tick size war in each of the Scandinavian markets. The statistic we are interested in is how often the small-tick venues improve on the primary exchange prices by either one (new) tick, or by multiple ticks. For each stock-day, we calculate the proportion of the day the small-tick venue quotes a price that:

1. Improves on the primary exchange price by any amount (in Figure 4, this corresponds to the sum of the x-axis ticks +1 to ≥ 5).
2. Improves on the primary exchange price by more than one tick (in Figure 4, this corresponds to the sum of the x-axis ticks +2 to ≥ 5).

We again focus on differences relative to Chi-X. Figure 5 plots these two series throughout the calendar year 2009 for the three Scandinavian markets. The *difference* between the two series in Figure 5 gives the fraction of time the best Chi-X price undercuts the primary exchange by a single tick. Across all three markets, immediately after its tick size reduction, Chi-X drastically increases the fraction of time for which it improves on the primary exchange price by at least one new tick (the red line) to around 60%. The fraction of time for which Chi-X provides price competition of more than one new tick (the blue line) jumps by much less on the introduction of the finer pricing grid, reaching a maximum of only 20%-30% of the trading day. The stark difference in these observed changes is almost immediately reversed to pre-war levels once the tick sizes are harmonized across trading venues.

Fig. 5: Quote placement at Chi-X relative to the primary exchange

The figures show the fraction of the day that Chi-X offers prices that (1) improve on the primary exchange price (red line) or (2) improve on the primary exchange by more than one tick (blue line). The exact measurement procedure is detailed in Section 5.1. We only consider stocks in the blue-chip indexes at COP, OSE, and STO.

Panel A: Denmark (Copenhagen)

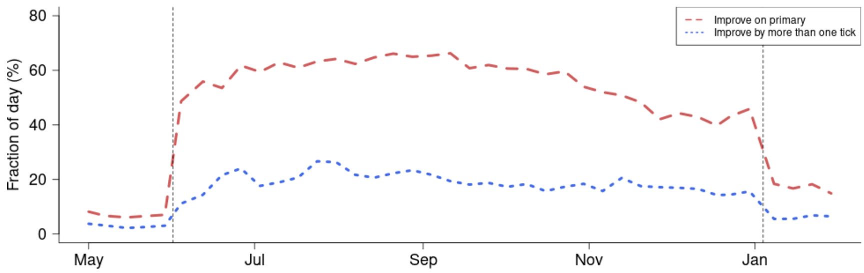

Panel B: Norway (Oslo)

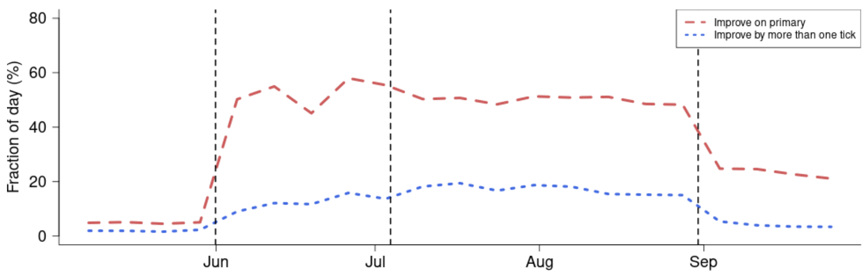

Panel C: Sweden (Stockholm)

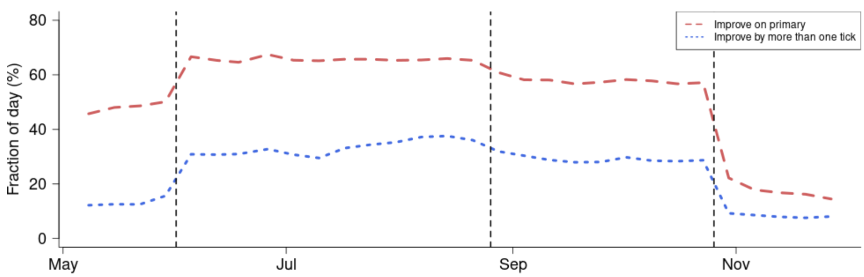

5.3. Mechanism: What drives price competition?

The time-series evidence in Figure 5 shows that market participants primarily used the small-tick trading venues for the purpose of single-tick undercutting, but did engage in more meaningful price competition almost 20% of the time. These graphs present market-wide averages for blue-chip index stocks only, and therefore may mask any heterogeneity generated by the substantial cross-sectional variation in the extent to which the tick size constrains the bid-ask spread. The results of our market quality analysis in Section 4 indicate that the competitive benefits of narrower tick sizes were concentrated in constrained securities. We now examine the cross-sectional differences in stock-level tick constraints to help us further understand the drivers of the quoting behavior documented earlier in this section.

Intuitively, suppose stocks have a “natural” bid-ask spread if the stocks were not constrained by the tick grid. If this “natural” spread is smaller than one (large) tick, we classify these stocks as constrained, and expect traders to migrate towards the small-tick venue to engage in price competition, thus increasing trading activity and order book depth on the alternative venue. If, however, the “natural” spread is larger than one tick, the small-tick venue will primarily be used for undercutting (as in Biais et al. 2010). If the mechanism that drives traders to either undercut or compete on the new finer pricing grid is, in fact, the extent to which the security’s bid-ask spread was constrained by the existing tick size, we would expect to see more price competition (and depth improvement) on the alternative venue when quotes are constrained on the primary exchange.

To test these predictions, we consider three variables as proxies for the degree of economically meaningful price improvement:

(a) The fraction of time Chi-X improves on the primary exchange by > 1 tick.
(b) The fraction of time Chi-X improves on the primary exchange by > 3 ticks.
(c) Order book depth at the best Chi-X prices.[26]

The first two variables capture the extent of between-venue price competition—the likelihood of moving beyond one (or three) ticks relative to the primary exchange. The third looks at the depth quoted on Chi-X. We test whether the extent to which the primary exchange is spread-constrained affects these variables, controlling for other

[26] Note that the Chi-X depth is measured only at the best prices at Chi-X; it is not the depth measure where we aggregate depth at prices equal to and improving on the main market quotes. As such, it serves as a lower bound of the price improving depth quoted on the small tick exchange.

factors. Specifically, we estimate the following cross-sectional regression model:

$$y_{it} = \alpha_0 + \beta \text{Constrained}_{it} + \gamma X_{it} + \varepsilon_{it}, \tag{6}$$

where y_{it} is the variable of interest—either (a), (b), or (c), as defined above—for stock i on date t. The key explanatory variable is Constrained_{it}, which measures whether trading at the primary exchange is constrained by the tick size. We consider three versions of Constrained_{it}. The first two are those introduced in Section 4.3: the binary and Kwan et al. (2015) constraint measures. Third, we follow O'Hara et al. (2018) and measure constraints by the fraction of the day quotes at the primary exchange are constrained to one tick. Unlike the first two measures, which are calculated based on pre-tick size war data, we calculate the third on a daily basis, and use the contemporaneous measure in our regressions. Our regressions control for the stock price and market capitalization (X_{it}).

Table 8. Are constrained stocks more likely to have Chi-X price competition?

The table reports estimates from regressions of the following form: $y_{it} = \alpha_0 + \beta \text{Constrained}_{it} + \gamma X_{it} + \varepsilon_{it}$, where y_{it} is either (a) the fraction of the trading day Chi-X improves on the primary exchange price by more than one tick (columns (1)-(3)), (b) the fraction of the trading day Chi-X improves on the primary exchange by more than three ticks (columns (4)-(6)), or (c) the (log) order book depth at the best Chi-X prices (columns (7)-(9)). X_{it} is a vector of control variables, which includes the (log) stock price and (log) market capitalization. The table only reports estimates of β. We consider three different measurements of Constrained_{it}. *Constrained(binary)* is an indicator variable that equals one if the stock is constrained according to the binary measure introduced in Section 4.3. *At-margin(Kwan et al.)* and *Constrained(Kwan et al.)* are indicator variables that equal one if the stock is At-margin or Constrained according to the Kwan et al. (2015) constraint measure introduced in Section 4.3. *Fraction at minimum tick* is the fraction of the trading day that the quotes on the primary exchange are constrained at one tick. All regressions use data from June 2009, the first month of the tick size war, and the estimation sample comprises all stock-days with trading at both Chi-X and the primary exchange. We only consider stocks in the blue-chip indexes at COP, OSE and STO. t-statistic in parentheses. The standard errors are adjusted for clustering. Significance levels are indicated as: ** p<10%*, *** p<5%*, and **** p<1%*.

	Dependent Variable								
	(a) Fraction > 1 tick			(b) Fraction > 3 ticks			(c) Chi-X Depth		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Constrained (Binary)	0.081***			0.066***			0.584***		
	(2.923)			(2.884)			(5.426)		
At-margin (Kwan et al.)		0.020			0.016			0.324***	
		(1.057)			(0.905)			(2.943)	
Constrained (Kwan et al.)		0.133***			0.114***			0.652***	
		(4.493)			(4.356)			(6.156)	
Fraction at minimum tick			0.235***			0.201***			1.240***
			(4.291)			(4.023)			(10.218)
N	1,393	1,393	1,393	1,393	1,393	1,393	1,393	1,393	1,393
$\overline{R}^2$	0.080	0.152	0.146	0.080	0.152	0.142	0.625	0.601	0.696

Table 8 presents estimates of β from Equation (6). The estimates are based on data from June 2009, the first month of the tick size war. The results show that the extent to which the primary exchange bid-ask spread is constrained by the tick size has a statistically significant and positive impact on both meaningful price competition and depth quoted at the alternative venue. Stocks whose bid-ask spreads are constrained by the tick size at the primary exchange are more likely to see price improvement beyond one tick, and have more depth quoted on Chi-X. The results are robust to all three measurements of Constrained_{it}. These results are consistent with the findings of Foucault and Menkveld (2008) that increased competition—in this case facilitated by a finer pricing grid—can enhance liquidity provision and, at the same time, increase order book depth.

The results presented in Table 8 are also consistent with the market quality improvements being concentrated in the constrained securities, as documented in Section 4. When the tick size itself prevents traders from quoting narrower spreads, market quality is likely to be hindered. The addition of a finer pricing grid on an alternate exchange facilitates market makers to explicitly compete against each other, reducing the costs of trading. These benefits are concentrated in the securities for which the level of constraint imposed by the existing tick size was the largest. This indicates that the benefits of finer pricing grids may be limited to situations where spreads are constrained by the pricing grid. Luckily, such a measure is easy to obtain, and holds promise as a tool for both regulators and stock exchanges to identify securities which will benefit the most from a revised pricing grid.

6. Conclusion

This paper empirically assesses the effects of between-venue competition on the tick size. We consider the European tick size wars in June, 2009, in which three new entrant trading venues unexpectedly reduced their tick sizes for stocks with primary listings at the Copenhagen, Oslo, and Stockholm exchanges. We find that the tick-size-reducing venues immediately attracted market shares of both executed and quoted volume from the large-tick primary exchanges. We also find that the competitive tick size reductions significantly improved measures of market quality, such as market-wide spreads, aggregate volume levels, and the informational efficiency of prices. The market quality gains from tick size competition were strongest in already-liquid stocks, where the spread was constrained by the existing tick size. Overall, the results suggest positive effects of competitive tick size reductions on standard measures of market quality.

We also shed light on the mechanisms through which the small-tick venues captured order flow from the primary exchanges. Our results suggest that, predominantly, the small-tick venues were used by liquidity providers to undercut the wider tick size in the primary exchanges by very small amounts—one "new" tick. Such undercutting is similar to what is observed with fractional price improvement in dark pools and inverted maker-taker venues in the United States and other jurisdictions (e.g., Comerton-Forde et al. (2019); Kwan et al. (2015)). However, in cases where stocks' trading costs were held artificially high by the tick size—approximately 20% of securities—we find that the introduction of a small-tick venue facilitated more meaningful price improvement—improving on the primary exchange spread by five or more "new" ticks. Although there are similarities between "synthetic" mechanisms to avoid a harmonized tick size schedule, such as dark trading or inverted maker-taker fees, these mechanisms do not provide participants with a new pricing grid on which to "compete" with each other, preventing a new, narrower equilibrium spread from emerging. This effectively precludes these market-structure changes from providing a beneficial "competitive" outcome, and rather enforces the sub-optimal "undercutting" equilibrium. It is in the stocks with the most cross-venue price competition—constrained stocks—we find the greatest market quality gains from competitive tick size reductions.

The importance of tick sizes in terms of encouraging liquidity in modern equity markets is currently the focus of considerable policy debate, particularly in light of the recent SEC tick size pilot in the United States and new regulations imposed by MiFID II in Europe to introduce harmonized tick size schedules. However, the effect of tick size competition has remained elusive as a research area, not least owing to the regulations in both in the United States and Europe that enforce a common tick size across trading venues. It is possible that the introduction of dynamic, non-binding tick size schedules that accurately reflect the cost of liquidity provision in markets could reduce complexity and segmentation in what have become increasingly complex and fragmented global equity markets. Our results show that the competitive reductions in tick sizes were beneficial for the overall marketplace and, in contrast to other mechanisms that aim to achieve the same outcome in a fixed tick-size environment, may lead to fewer distortions in the allocation of wealth, segmentation of traders, and leakage of information that can occur between market participants.

References

Michael Aitken, Haoming Chen, and Sean Foley. The impact of fragmentation, exchange fees and liquidity provision on market quality. *Journal of Empirical Finance*, 41:140–160, March 2017. doi: 10.1016/j.jempfin.2016.10.002.

Rui Albuquerque, Shiyun Song, and Chen Yao. The price effects of liquidity shocks: A study of the sec's tick size experiment. *Journal of Financial Economics*, 2020. doi: 10.1016/j.jfineco.2020.07.002.

James J Angel, Lawrence Harris, and Chester S Spatt. Equity trading in the 21st century. *Quarterly Journal of Finance*, 1(1):1–53, 2011.

V Ravi Anshuman and Avner Kalay. Market making with discrete prices. *The Review of Financial Studies*, 11(1):81–109, 1998.

Autorité des marchée financiers (AMF). MiFID II: Impact of the new tick size regime. Working Paper, AMF, available at www.amf-france.org, March 2018.

Robert P Bartlett, III and Justin McCrary. Subsidizing liquidity with wider ticks: Evidence from the tick size pilot study. *Journal of Empirical Legal Studies*, 17(2):262–316, 2020. doi: 10.1111/jels.12252.

BATS. Pan European Tick Size Pilot. BATS Trading Working Paper, July 2009.

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

Hendrik Bessembinder. Trade execution costs and market quality after decimalization. *Journal of Financial and Quantitative Analysis*, pages 747–777, 2003.

Ayan Bhattacharya and Gideon Saar. Limit order market under asymmetric information. Working Paper, 2020.

Bruno Biais, Christophe Bisière, and Chester Spatt. Imperfect competition in financial markets: An empirical study of Island and Nasdaq. *Management Science*, 56(12):2237–2250, 2010.

Jonathan Brogaard, Bjorn Hagströmer, Lars Nordén, and Ryan Riordan. Trading fast and slow: Colocation and liquidity. *Review of Financial Studies*, 28(12):3407–3427, December 2015.

Sugato Chakravarty, Robert A. Wood, and Robert A. Van Ness. Decimals and liquidity: A study of the NYSE. *Journal of Financial Research*, 27(1):75–94, 2004. doi: 10.1111/j.1475-6803.2004.00078.x.

Yong Chao, Chen Yao, and Mao Ye. Why Discrete Price Fragments U.S. Stock Exchanges and Disperses Their Fee Structures. *The Review of Financial Studies*, 32(3):1068–1101, 07 2018. doi: 10.1093/rfs/hhy073.

Daniel Chen and Darrell Duffie. Market Fragmentation. *American Economic Review*, 111 (7):2247–2274, July 2021. doi: 10.1257/aer.20200829.

Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Evidence on the speed of convergence to market efficiency. *Journal of Financial Economics*, 76(2):271–292, 2005. doi: 10.1016/j.jfineco.2004.06.004.

Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Liquidity and market efficiency. *Journal of Financial Economics*, 87:249–268, 2008.

Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Recent trends in trading activity and market quality. *Journal of Financial Economics*, 101:243–263, 2011.

Kee H. Chung, Albert Lee, and Dominik Rösch. Tick size, liquidity for small and large orders, and price informativeness: Evidence from the tick size pilot program. *Journal of Financial Economics*, 136(3):879–899, 2020. doi: 10.2139/ssrn.3220470.

Jean-Edouard Colliard and Thierry Foucault. Trading fees and efficiency in limit order markets. *Review of Financial Studies*, 25:3389–3421, 2012.

Carole Comerton-Forde, Vincent Grégoire, and Zhuo Zhong. Inverted fee structures, tick size, and market quality. *Journal of Financial Economics*, 134(1):141 – 164, 2019. doi: 10.1016/j.jfineco.2019.03.005.

Jennifer Conrad and Sunil Wahal. The term structure of liquidity provision. *Journal of Financial Economics*, 136(1):239–259, 2020.

Patrick J. Dennis and Patrik Sandås. Does trading anonymously enhance liquidity? *Journal of Financial and Quantitative Analysis*, 55(7):2372–2396, 2020.

Anne Dhyrberg, Sean Foley, and Jiri Svec. When bigger is better: The impact of a tiny tick size on undercutting behavior. Working paper, SSRN, June 2019.

Sean Foley and Tălis J Putniņš. Should we be afraid of the dark? Dark trading and market quality. *Journal of Financial Economics*, 122(3):456–481, 2016.

Sean Foley, Elvis Jarnecic, and Anqi Liu. Forming an orderly line - how queue-jumping drives excessive fragmentation. Working paper, SSRN, November 2019.

Thierry Foucault and Albert J Menkveld. Competition for order flow and smart order routing systems. *Journal of Finance*, 63(1):119–158, 2008.

Thierry Foucault, Ohad Kadan, and Eugene Kandel. Limit order book as a market for liquidity. *The Review of Financial Studies*, 18(4):1171, 2005.

Thierry Foucault, Ohad Kadan, and Eugene Kandel. Liquidity cycles and make/take fees in electronic markets. *Journal of Finance*, 68(1):299–341, 2013.

Ronald L Goettler, Christine A Parlour, and Uday Rajan. Equilibrium in a dynamic limit order market. *Journal of Finance*, 60(5):2149–2192, 2005.

Todd G Griffith and Brian S Roseman. Making cents of tick sizes: The effect of the 2016 U.S. SEC tick size pilot on limit order book liquidity. *Journal of Banking and Finance*, 101:104 – 121, 2019. doi: 10.1016/j.jbankfin.2019.01.017.

Peng William He, Elvis Jarnecic, and Yubo Liu. The determinants of alternative trading venue market share: Global evidence from the introduction of Chi-X. *Journal of Financial Markets*, 22:27–49, 2015.

Hans Hvide, Magne Mogstad, Tom Meling, and Ola Vestad. Broadband internet and the stock market investments of individual investors. Working paper, SSRN, January 2021.

Kjell Jørgensen, Johannes A Skjeltorp, and Bernt Arne Ødegaard. Throttling hyperactive robots - Order to trade ratios at the Oslo Stock Exchange. *Journal of Financial Markets*, 37(1):1–16, 2018.

Ohad Kadan. So who gains from a small tick size? *Journal of Financial Intermediation*, 15(1):32–66, 2006.

Jon Kerr, Gil Sadka, and Ronnie Sadka. Illiquidity and Price Informativeness. *Management Science*, 66(1):334–351, January 2020. doi: 10.1287/mnsc.2018.3154.

Amy Kwan, Ronald Masulis, and Thomas H McInish. Trading rules, competition for order flow and market fragmentation. *Journal of Financial Economics*, 28(2):592–636, 2015.

Andrew W Lo and A Craig MacKinlay. Stock market returns do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies*, 1:41–60, 1988.

Laura Malceniece, Kārlis Malcenieks, and Tālis J. Putniņš. High frequency trading and comovement in financial markets. *Journal of Financial Economics*, 134:381–399, 2019. doi: 10.1016/j.jfineco.2018.02.015.

Katya Malinova, Andreas Park, and Andriy Shkilko. Design report for the CSA pilot study on rebate prohibition (revised to address public comments). August 2019.

Tom Grimstvedt Meling. Anonymous trading in equities. *The Journal of Finance*, 76(2): 707–754, 2021. doi: 10.1111/jofi.12988.

Albert J Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16:712–740, 2013.

Albert J. Menkveld, Bart Zhou Yueshen, and Haoxiang Zhu. Shades of darkness: A pecking order of trading venues. *Journal of Financial Economics*, 124(3):503 – 534, 2017.

Randi Næs, Johannes A Skjeltorp, and Bernt Arne Ødegaard. Stock market liquidity and the Business Cycle. *Journal of Finance*, LXVI:139–176, February 2011.

Maureen O'Hara, Gideon Saar, and Zhuo Zhong. Relative tick size and the trading environment. *Review of Asset Pricing Studies*, 9(1):47–90, 12 2018. doi: 10.1093/rapstu/ray009.

Marco Pagano. Trading volume and asset liquidity. *Quarterly Journal of Economics*, 104 (2):255–274, May 1989.

Barbara Rindi and Ingrid M Werner. U.S. tick size pilot. Fisher College of Business Working Paper No. 2017-03-018, March 2019.

Johannes Skjeltorp and Bernt Arne Ødegaard. Why do listed firms pay for market making in their own stock? *Financial Management*, pages 241–261, Summer 2015.

Chester Spatt. The new realities of market structure and liqudity: Where have we been? Where are we going? In *Public Policy and Financial Economics: Papers in Honor of Professor George G. Kaufman.* World Scientific, 2017.

Vincent van Kervel and Albert J Menkveld. High-frequency trading around large institutional orders. *The Journal of Finance*, 74(3):1091–1137, 2019. doi: 10.1111/jofi.12759.

Chen Yao and Mao Ye. Why Trading Speed Matters: A Tale of Queue Rationing under Price Controls. *The Review of Financial Studies*, 31(6):2157–2183, 04 2018. doi: 10.1093/rfs/hhy002.

Acknowledgements

Earlier versions of the paper have circulated under titles *Tick Size Competition, High Frequency Trading, and Market Quality* and *Tick Size Wars: Competitive Tick Size Regimes and Trader Behavior*. This paper has benefited from discussions with Tamás Bátyi, Terrence Hendershott, Hans K. Hvide, Teis Lunde Lømo, Christine Parlour, Talis Putnins, Bjørn Sandvik, Eirik A. Strømland, and Jonas Tungodden. The paper has also benefited from the comments of two anonymous referees. We are grateful for the comments from seminar participants at the "Conference on the Econometrics of Financial Markets," the "2^{nd} Paris-Dauphine Workshop on Microstructure," the "2019 CEPR-Imperial Plato Market Innovator (M13) Conference," the "3rd SAFE Market Microstructure Conference." and the University of Technology, Sydney. We are particularly grateful for the comments from our discussants Gernot Doppelhofer, Carole Gresse, Stig Lundeby, Konrad Raff, and Christine Parlour.