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

Sean Foley Tom Grimstvedt Meling Bernt Arne Ødegaard*

August 2020
Preliminary

Abstract

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

Keywords: Equity Trading; Limit Order Markets; Tick Sizes
JEL Codes: G10; G20

*Foley is at Macquarie University. `sean.foley@mq.edu.au`. Meling is at the University of Chicago and the University of Bergen. `Tom.Meling@uib.no`. Ødegaard is at the University of Stavanger. `bernt.a.odegaard@uis.no`. Earlier version 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 László Bátyi, Terrence Hendershott, Hans K. Hvide, Teis Lunde Lømo, Christine Parlour, Talis Putnins, Bjørn Sandvik, Eirik A. Strømland, and Jonas Tungodden. We are grateful for 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 comments from our discussants Gernot Doppelhofer, Carole Gresse, Stig Lundeby, Konrad Raff, and Christine Parlour.

Introduction

Fragmentation is a key feature of modern equity markets. Spurred by regulatory initiatives to introduce competition to monopolistic primary exchanges, notably RegNMS 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 rules and "synthetically" compete on tick sizes through alternative trading mechanisms such as dark pools, batch auctions and inverted maker-taker venues. Understanding the connection between tick sizes, fragmentation and between-venue competition has generated a significant body of both empirical and theoretical work.

Empirically, the connection between equity market fragmentation and a desire to circumvent existing tick size regimes has been well documented in the context of inverted maker-taker venues in recent papers by e.g. Spatt (2017) and Comerton-Forde et al. (2019). The proliferation of dark pool venues in the U.S. 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 cost. 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 is 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". These arguments are 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) further 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. Such theoretical tensions make assessing the role of tick size competition on market quality and fragmentation an empirical question.

In this paper, we examine the effects of the European "tick size wars" where stock ex-

[1]Most recently the SEC 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).

changes 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 before 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, allowing us to understand how both liquidity suppliers and market participants respond to the dynamically evolving landscape, without needing to 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 on pricing grids.

Our main result is that the exchanges that reduced their minimum price increments during the tick size war immediately captured market shares of both quoted and executed volume from exchanges that kept their price increments wide, with liquidity suppliers particularly attracted to the small-tick venues. The existing ELPs 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. This competition improved aggregate market quality, we find, reducing transactions costs and increasing available depth. We find these improvements to be strongest in stocks where the bid-ask spread was constrained by the old, wider tick size. Our findings are in many ways 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 empirical 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 additional fragmentation. This is consistent with the theoretical arguments of Colliard and Foucalt (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. Our findings are 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 observed directly in our analysis via tick size competition, further connect to the empirical findings of papers such as Kwan et al. (2015), Foley and Putniņš (2016), and Menkveld et al. (2017) on the desirability of dark pools, as well as those of Menkveld (2013), Battalio et al. (2016) and Comerton-Forde et al. (2019) on why fee structures (in general) and inverted maker-taker venues (specifically) may be preferred to traditional venues. In our setting, traders are observed to utilize the option to undercut by a very small fraction in order 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 where securities are heavily constrained by the tick size. Showing that these effects are observed when exchanges are allowed to compete directly on pricing grids 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 those securities whose 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 tick sizes are heavily constrained, the introduction of a new, lower tick size regime results in market participants significantly driving down transactions costs and increasing quoted depth. This occurs as liquidity providers compete on the finer pricing grid, narrowing the spread by several new ticks. These results are consistent with the theoretical work of Anshuman and Kalay (1998) that discrete pricing grids generate rents for market makers, as well as the results of Goettler et al. (2005), that lower tick sizes will 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 U.S. primarily benefited the most liquid securities.

Our third main contribution is to document the effects of tick size competition on trader behavior, particularly for ELPs. We find evidence consistent with a "separating equilibrium", where ELPs "undercut" the existing pricing grid by either one or many new ticks. When the ELPs are provided with a finer pricing grid, they almost always undercut existing prices by the minimum amount possible. Competition among the ELPs forces bid-ask spreads to decrease in the stocks which were previously the most heavily constrained, consistent with the theoretical predictions of Kadan (2006) of more aggressive competition between market makers in the presence of reduced tick sizes. This indicates that identifying and reducing tick sizes for such securities can greatly improve measures of market quality.

Our findings also contribute to the existing policy debate on optimal tick size regimes. The results of our paper 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, our findings suggest. Moreover, recent analyses of the SEC's now ended Tick Size Pilot provide conflicting evidence on the market quality effects of tick size changes.[2] The inability of this literature to

[2]Chung et al. (2020) show that increased tick sizes have resulted in costs for small orders but benefits for larger orders, reducing 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

reach a consensus on tick size regimes, combined with the intentions of MiFiD II to reform tick sizes[3], speaks to the importance for regulators of additional evidence, such as ours, on how changes to tick size regimes impact trader behavior and welfare.

More broadly, our findings also connect to policy debates on optimal maker-taker regimes, as exemplified by the SEC and Canadian Securities Administration current plans for pilots to remove maker-taker fees (see 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 for example 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 1 provides institutional details and summarizes the main event of the "Tick Size Wars". Section 2 provides details on the data used in the empirical analysis. Section 3 presents our main empirical results. Section 4 analyzes changes to order submission strategies during the tick size war. Finally, Section 5 concludes.

1 Institutional details

By the early 2000's, the European national stock exchanges, such as those in London, Copenhagen, Oslo and Stockholm, had all finished the transition to pure electronic limit order markets. At the time, they held dominant positions in their respective markets. This dominance was broken by the MiFID reform in 2007. MiFID unleashed competition for European order flow, causing equity trading to fragment across a wide range of trading venues.

1.1 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 using the same tick size schedules. The MTFs started by trading the largest

benefits of the Tick Size Pilot have currently been ascribed primarily 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.

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

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 — i.e., the most liquid stocks — while 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 quickly christened the "Tick Size War." The MTFs maintained lower 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]

1.2 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 — COP, OSE, and STO — by reducing the tick size for their 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 COP, OSE, or STO primary listing. Turquoise followed on June 8 by reducing its tick sizes for all Scandinavian stocks and five LSE listings. Finally, on June 15, BATS reduced its tick sizes for all Scandinavian stocks, as well as ten stocks on the LSE and five listed in Milano (BATS, 2009).

1.3 Tick sizes: The battleground

The MTF tick size reductions during the *break-out phase* were substantial. Take the OSE as an example. At the time of Chi-X's June 1 tick size reduction, the OSE operated with three tick size schedules, all with ticks varying between NOK 0.01 and 1, where the price cutoffs varied depending on index membership (i.e. 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,

[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, while van Kervel and Menkveld (2019) examines the impact of HFT on institutional trading costs. Using Norwegian data, Meling (2020) estimates the effects of post-trade anonymous trading, while Næs et al. (2011) explore the connection between market liquidity and the business cycle.

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.

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 reduced its tick size uniformly to 0.01 for trading in their 25 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 ten highly liquid stocks, allowing between-venue tick size differences to persist for less liquid stocks.

1.4 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

[6]Section A1 of our Internet Appendix provides more detailed summaries of the tick size changes.

[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.

detrimental effects for 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, while in Sweden and Denmark the primary exchange tick size decreased, while those of the competing venues increased.

2 Data

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

2.1 Data sources

We obtain trade-and-quote data from the Thomson Reuters Tick History (TRTH) database. For trading venues with displayed limit order books, the TRTH provides data on all executed trades, as well as the ten best levels of the order book. We collect data for the trading in all Scandinavian securities on their primary exchanges as well as from the trading of these securities on Chi-X, BATS and Turquoise. We also collect data on the number of shares outstanding and index constituency in each market's blue-chip index for the period of 2008–2010. In the London market, the tick size war lasted less than two weeks; in the main analysis, we focus on the longer-lived events in the Scandinavian markets.[8]

We use the TRTH trade-and-quote data to construct several standard measures of market quality, including the quoted, effective and realized bid-ask spreads, the price impact, as well as measures of order book depth and realized volatility. The construction of these market quality measures follow standard procedures in the market microstructure literature, and are conducted at the 1 minute horizon where a time is required (i.e. 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 specific definitions of each of the measures can be found in Section A2.1 of the Internet Appendix.

[8]In the Internet Appendix, we provide empirical results for LSE listed securities.

2.2 Measuring aggregate market quality

In addition to the standard venue-specific measures of market quality, we construct measures of aggregate market quality across all trading venues in a given market. While the European markets have no formal notion of a national best bid and offer (NBBO),[9] we find it useful to construct an estimate of the best bid and offer prices across the competing venues. We do so by aligning contemporaneous order books across venues and, at each point in time with an update to the order book, record the best bid and ask prices. For simplicity we term these estimates the "NBBO." These NBBO prices are then used to construct market-wide estimates of transactions costs, which are based on the NBBO instead of the venue-specific bid and ask prices.

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 which still improve on the large-tick venue. In our view, the most relevant comparison is to aggregate all depth quoted on the small-tick venue at prices which are equal to (or better than) the best prices at the large-tick exchange, as illustrated by the shaded area in 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.

2.3 Descriptive statistics

Panel A of Table 2 summarizes 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. While 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, in each market covering significant numbers of index and non-index stocks. Our empirical 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 the competing MTFs. Table 2 suggest that market quality, as measured by spreads, order book depth, and trading volume, is significantly better at the primary exchanges than

[9] 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.

Table 2: Summary statistics

Panel A summarizes 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 COP, 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 uses 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 ten-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 crossection				
		Primary exchange		Chi-X	BATS	Turquoise
Country	Index	in index	outside 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 (thous)	2240	1016	7686	187	69	59815	488	260	13870
Kroner Volume (mill)	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
Order to Trade	11.8	6.8	7686	16.8	7.3	62841	282.2	52.8	14830
Market Cap (mill)	50748	21440		3242	406				

Figure 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 depth at the best bid and ask on that exchange. That is compared with the sum of depth at prices equal to or improving on the bid and ask prices on the small-tick exchange. The shaded areas indicate the comparison.

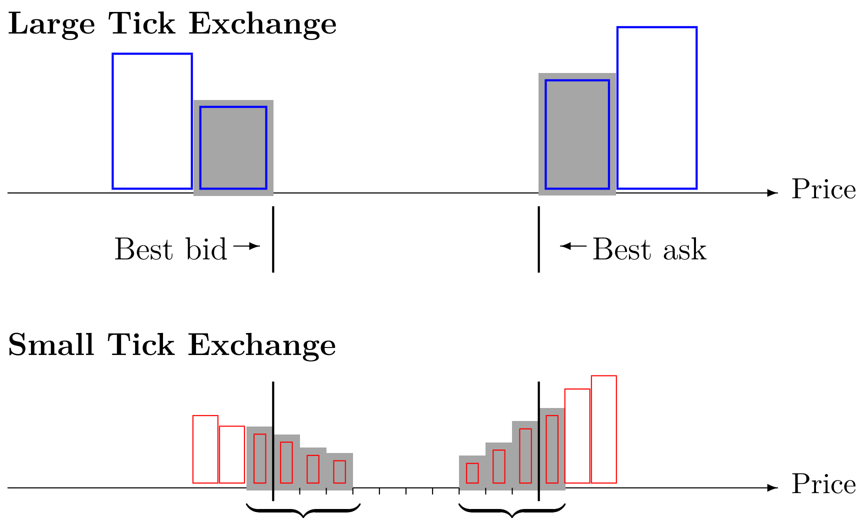

the MTFs. For example, the average quoted and effective bid-ask spreads on the MTFs are 1% and 0.32%, about three times larger than the average on primary exchanges. In terms of depth, the MTFs average about 480 thousand kroner — about a fifth of the primary exchange average of 2.2 million. In terms 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 to 0.69% on the primary exchanges — perhaps due to the less frequent trading.

Finally, we summarize our market-wide spread measures. As adding the competing MTFs can only improve the market-wide quoted spread, the average NBBO quoted spread (0.32%), is lower than the 0.36% average for the primary exchange, 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 NBBO numbers for the MTFs are slightly different from those of the primary exchange stocks, due to differences in the coverage of stocks.

3 Main results

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

3.1 Did the tick-size reducing MTFs capture market share?

We consider two measures of market shares. The first, a post-trade measure, is simply a 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.[10] We examine both pre and post-trade market shares in order 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: COP, OSE, and COP. The graphical evidence is striking. Immediately after the start of the tick size war, on June 1, 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% and 3.5%.[11] 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%.

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 the Stockholm market, the corresponding effects are somewhat less extreme, possibly reflecting the fact that Chi-X already in the pre-war period offered smaller tick sizes than 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

[10] 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.

[11] These numbers represent estimates from a before-after regression specification. The regression estimates are presented in Internet Appendix Section A3.1.

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

Our graphical results presented in Figure 2 provide compelling evidence that competing trading venues can capture market shares, both in terms of quoted and executed volume, from incumbent stock exchanges by competing directly on the minimum price increment. While 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, while 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".

3.2 Did the tick size war affect market quality?

Having documented significant effects of the tick size war on both pre-trade and post-trade market shares, we next consider its effects on market quality, both at individual exchanges and in the aggregate. As a starting point, we estimate these effects using a simple before-after specification surrounding the start of the tick size war:

$$y_{it} = a + \beta Post_{it} + \gamma \mathbf{X}_{it} + \varepsilon_{it}, \tag{1}$$

where y_{it} is the outcome of interest, e.g. the effective spread, for stock i on date t;[12] $Post_{it}$ is an indicator variable that equals zero before the start of the tick size war and one after; $\mathbf{X}_{it}$ is a vector of stock-level control variables; and ε_{it} is an error term. The coefficient of interest is β, which captures the change in y_{it} from before to after the start of the tick size war. We estimate β using a short sample period (May 15 – June 30) to exclude the primary exchanges' tick size retaliations, and include in the estimation sample all stocks that were traded on both the primary and MTF venues before the start of the tick size war.

Pooling data across venues, Table 3 presents one estimate of β for the primary exchanges (COP, OSE and STO) and one estimate of β for the alternative venues (Chi-X, BATS and

[12]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 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 by traded venue. When considering aggregate measures of market quality, such as the NBBO, Statoil is represented by the same NBBO for both primary and MTFs venues.

Figure 2: Effects of tick size war on market shares

The figure plots two measures of primary exchange market shares throughout the tick size war. In black, as our pre-trade market share measure, we plot the fraction of time the primary exchange quotes at the market-wide best bid and ask prices. In 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 1 June 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)

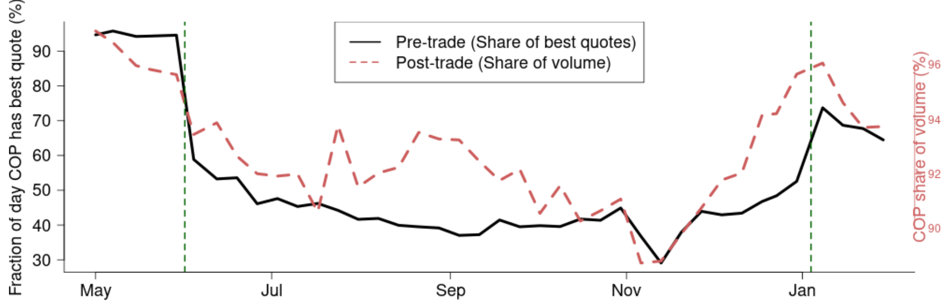

Panel B: Norway (Oslo)

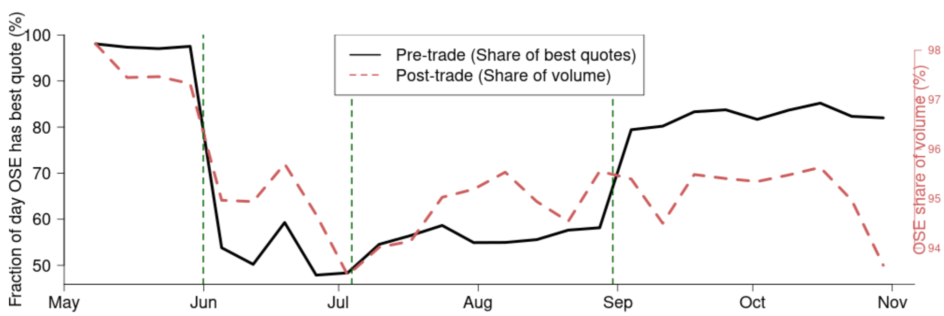

Panel C: Sweden (Stockholm)

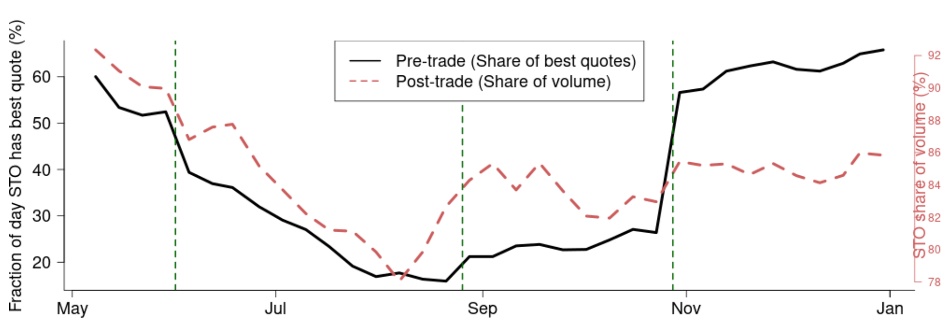

Table 3: Before-after effects of the tick size war on market quality

The table presents estimates of β from the following before-after regression specification: $y_{it} = a + \beta Post_{it} + \gamma \mathbf{X}_{it} + \varepsilon_{it}$. $Post_{it} = 1$ for all time periods $t \geq t^*$, where t^* represents the start of the tick size war for a given trading venue. t^* = June 1, 2009 for OSE, COP, STO, and CHI. t^* = June 8, 2009 and June 15, 2009 for Turquoise and BATS, respectively. $\mathbf{X}_{it}$ is a vector of control variables, which includes the natural logarithm of the stock price and stock-level fixed effects. The sample period is May 15 to June 30. The estimation sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The coefficient β is estimated separately for the primary exchanges, the MTFs, and the NBBO (Aggregate). *Quoted*, *Effective* and *Realized* spreads, and *Price Impact* are in basis points. *Depth* is order book depth, transformed with the natural logarithm. *Volatility* is measured in percentage points. *Volume* is the trading volume, transformed with the natural logarithm. For ease of exposition, we have included N, the number of observations, only for regressions using *Volume* as the dependent variable. N is very similar across outcomes. t-statistics are presented in parentheses. Standard errors are clustered at the RIC-level. Significance levels indicated as: ** p<10%, ** p<5%, *** p<1%*

	Primary	MTFs	Aggregate
Quoted spread	−0.18	−18.97***	−2.64***
	(−0.20)	(−6.31)	(−2.72)
Effective spread	−0.38	−4.67***	−0.14
	(−1.22)	(−4.80)	(−0.28)
Realized spread	−0.35	−1.48*	0.94*
	(−1.08)	(−1.72)	(1.67)
Price impact	−0.04	−3.10***	−0.39
	(−0.17)	(−2.94)	(−1.21)
Depth	−0.02	−0.05**	0.08***
	(−0.80)	(−2.35)	(3.14)
Volatility	−0.12	−0.01***	−0.13
	(−1.02)	(−3.19)	(−1.16)
Volume	−0.13***	0.24***	−0.10***
	(−3.48)	(5.30)	(−2.66)
N	2697	5886	2697

Turquoise). For the primary exchanges, where the tick size did not change in June 2009, we find a statistically significant volume reduction of about 13%.[13] For the average stock trading on the primary exchange, we find no statistically or economically significant effects of the tick size war on measures of spreads, order book depths, or realized volatility. For the alternative venues, where tick sizes were reduced during June 2009, we find a considerable increase in trading volume (+20%) accompanied by reductions in both quoted spreads ($\beta = -19$ basis points) and effective spreads ($\beta = -4.67$ bp). The latter effect is driven primarily by a reduction in the price impact of market orders. We also find minor reductions in alternative venue depth (−5%) and return volatility (−0.05 percentage points).

What do the individual-venue effects in Table 3 imply for aggregate market quality? The observed improvements in market quality on the alternative venues may not necessarily affect aggregate market quality in a meaningful way, given that only a small share of overall trading takes place there (less than 10%, as shown in Figure 2). To empirically assess the aggregate market quality effects of the tick size war, we re-estimate eq. (1) using measures of overall market quality based on NBBO calculations, together with total trading volume across all venues. The results are presented in the "Aggregate" columns of Table 3, and show that the decrease in alternative venue quoted spreads generated a significant decrease in the NBBO quoted spread ($\beta = -2.64$ bps). In terms of volumes, we find that the sizable increase in volume on alternative venues (+24%) was more than offset by the reduction in primary exchange volume (−10%), causing a reduction in aggregate trading volume of about 10%. We also find that market-wide depth increased by more than 8%.

Overall, the before-after estimates in Table 3 show that the tick size war of June 2009 led to a considerable shift in volume *away* from primary exchanges *towards* the small-tick size alternative venues, accompanied by significant reductions in spreads at the alternative venues, yielding lower market-wide transactions costs.

We further complement this analysis with a more robust econometric specification. The before-after estimates may be biased by market-wide trends unrelated to the tick size war. For example, trading activity and liquidity may be systematically lower during June (the post period) than May (the pre period), due to the reduced trading historically observed during the Scandinavian summer (holiday) months. To address this concern, we follow a comprehensive literature in economics (e.g. Bertrand et al. 2004) and finance (e.g. Colliard and Hoffmann 2017) and use a control group of stocks unaffected by the tick size war to

[13] Specifically, we find that the natural logarithm of trading volume is reduced by 0.13, which corresponds approximately to a 13% reduction in the level of trading volume

estimate a difference-in-differences model that adjusts for the impact of market-wide trends on the estimated effects.

Our difference-in-differences approach exploits the fact that, by June 2009, only the most liquid stocks in the Scandinavian markets could be traded outside the primary exchanges and, for this reason, only this subgroup of stocks was affected by the alternative venues' tick size reductions. All other stocks could only be traded on their primary exchange and were therefore *not* affected by the tick size war. Accordingly, our "treated" group is comprised of all stocks traded on alternative venues prior to June 2009 — i.e., the same stocks as in Table 3. As our "control" group, we use all primary exchange stocks that were *not* traded on any of the alternative venues throughout 2009. In the differences-in-differences design, we compare the evolution of market quality for "treated" and "control" stocks from before to after the start of the tick size war, using the following difference-in-differences model:

$$y_{it} = \alpha_i + \alpha_t \times \alpha_c + \tau D_{it} + \omega_{it}, \tag{2}$$

where $D_{it} = Post_{it} \times Treated_{it}$ 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 0 otherwise; α_i are stock-level fixed effects; 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 equation (2) separately for primary exchanges and alternative venues, and for brevity only report estimates of the average difference-in-differences effect, τ.

The identifying assumption of the difference-in-differences model is that treated and control group stocks follow parallel trends in their y_{it} absent the treatment of interest — i.e., the tick size war. In Figure A10 of the Internet Appendix we summarize key characteristics of our treated and control stocks. A main take-away is that the treated and control stocks in our sample differ widely in the levels of their y_{it}'s: for example, the average bid-ask spread in the control group is more than ten times wider than in the treated group. A natural concern is that a relatively small change to the y_{it}'s of the control group may overwhelm any changes to the y_{it}'s of the treated group, constituting a violation of the parallel trends assumption. To minimize this concern, we transform all our main market quality outcomes with the natural logarithm and, as a consequence, restate the parallel trends assumption in relative rather than absolute terms — with logarithms, we need the *percentage* change in y_{it} to be parallel between treated and control stocks. Figure A11 of the Internet Appendix shows, using data from before the tick size war, that treated and control stocks indeed follow remarkably similar trends in their log market quality characteristics, thus satisfying the parallel trends assumption of the difference-in-differences design.

Table 4: DiD effects of the tick size war on market quality

The table presents estimates of τ from the following difference-in-differences specification: $y_{it} = \alpha_i + \alpha_t \times \alpha_m + \tau D_{it} + \mathbf{X}_{it} + \omega_{it}$. $Post_{it} = 1$ for all time periods $t \geq t^*$, where t^* represents the start of the tick size war for a given trading venue. $t^* =$ June 1, 2009 for COP, OSE, STO, and CHI. $t^* =$ June 8, 2009 and June 15, 2009 for Turquoise and BATS, respectively. $D_{it} = Post_{it} \times Treated_{it}$, where $Treated_{it}$ is an indicator variable equal to one for stocks in the treated group. α_i, α_m, and α_t represent RIC, market (e.g., Norway), and date-level fixed effects, respectively. $\mathbf{X}_{it}$ is a vector of control variables, which includes the log of the stock price. The sample period is May 1 to June 30. The *treated* sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The *control* sample comprises stocks that could not be traded outside the primary exchanges throughout the calendar year 2009. In total, there are 89 treated units and 577 control units. The difference-in-differences coefficient τ is estimated separately for the primary exchanges, the MTFs, and the NBBO (Aggregate). All outcome variables except Volatility, which is measured in percentage points, have been transformed with the natural logarithm. For ease of exposition, we have included N, the number of observations, only for regressions using *Volume* as the dependent variable. N is very similar across outcomes. t-statistics are presented in parentheses. Standard errors are clustered at the RIC-level. Significance levels indicated as: ** p<10%, ** p<5%, *** p<1%*

	Primary	MTFs	Aggregate
Quoted spread	−0.08***	−0.33***	−0.20***
	(−4.40)	(−12.30)	(−8.39)
Effective spread	−0.08***	−0.27***	−0.05***
	(−3.99)	(−10.73)	(−2.62)
Realized spread	−0.15***	−0.31***	−0.04
	(−3.93)	(−7.29)	(−1.22)
Price impact	−0.04	−0.24***	−0.18***
	(−1.40)	(−5.66)	(−4.59)
Depth	0.00	−0.00	0.11***
	(0.10)	(−0.16)	(3.58)
Volatility	0.01	0.03	−0.01
	(0.10)	(1.21)	(−0.05)
Volume	0.12***	0.66***	0.16***
	(2.65)	(13.68)	(3.40)
N	23344		

The difference-in-differences estimates are presented in Table 4. 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. These positive effects of the tick size war on primary exchange market quality differ markedly from those documented in Table 3 — i.e., reduced volume, no effect on spreads — which may be consistent with the before-after estimates mistakenly attributing reduced summer-time market activity as an effect of the tick size war. For the alternative venues, the results are qualitatively similar to those for the simple before-after estimates in Table 3, though for the most part quantitatively more significant. For the aggregate market, the difference-in-differences specification shows a significant reduction in bid-ask spreads, and significant increases in both quoted depth and traded volume, confirming that the tick size war had beneficial effects for overall market quality.

3.3 Do binding tick sizes matter?

We have seen that the tick size war during June 2009 caused a reduction in aggregate trading costs and an increase in aggregate trading volume. To better understand the economic 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 best quotes, and reducing trading costs for all market participants.

On the other hand, 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.[14] Chung et al. (2020) find that a market-wide tick size *increase* results in costs for small orders, while accruing 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.

[14]Such undercutting behavior, experienced in an extremely small tick size environment for cryptocurrencies, has been shown to increase measures of transactions costs, discouraging limit order provision (Dhyrberg et al., 2019).

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.[15] Specifically, for stock i on date t, we start by calculating:

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

As $\text{TickPerSpread}_{\text{it}}$ becomes smaller, the quoted spread comes closer to being constrained by the tick size. In the limit, where $\text{TickPerSpread}_{\text{it}} = 1$, the tick size fully *constrains* the quoted bid-ask spread, preventing it from becoming smaller. Then, we base a binary categorization on the average of $\text{TickPerSpread}_{\text{it}}$ on the primary exchange one month prior to the tick size war. Stocks with $\text{Mean}(\text{TickPerSpread}_{\text{it}}) \geq 2$ are defined as "Unconstrained", while stocks with $\text{Mean}(\text{TickPerSpread}_{\text{it}}) < 2$ as "Constrained".

In Table 5 we re-estimate the difference-in-differences effect τ separately for treated stocks classified as *Constrained* and *Unconstrained*. We find that most of the observed market quality effects are driven by the sample of *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, while for *Unconstrained* stocks we only find a significant increase in trading volume on the alternative venues (+75%).

In columns labeled "Aggregate", Table 3 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 to 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 the securities whose

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

Table 5: The market quality effects of the tick size war for constrained and unconstrained stocks — Binary constraint measure

The table reports estimates of τ from the following difference-in-differences specification (the same as in Table 4): $y_{it} = \alpha_i + \alpha_t \times \alpha_m + \tau D_{it} + \mathbf{X}_{it} + \omega_{it}$, where α_i, α_m, and α_t represent RIC, market (e.g., Norway), and date-level fixed effects, respectively. $D_{it} = 1$ for all treated group observations on dates $t \geq t^*$, where t^* is the start date of the tick size war for a given trading venue. $\mathbf{X}_{it}$ is a vector of control variables, which includes the natural logarithm of the stock price. The difference-in-differences coefficient τ is estimated separately for the primary exchanges, the MTFs, and the NBBO (Aggregate). The baseline *treated* sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The *control* sample comprises stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. We split the *treated* sample into *Constrained* and *Unconstrained* stocks, as defined by the binary constraint measure described in Section 3.3, and estimate the difference-in-differences model separately for each of the treated samples, holding fixed the definition of the control sample. All outcome variables except *Volatility*, which is measured in percentage points, have been transformed with the natural logarithm. For ease of exposition, we present N, the number of observations, only for regressions using *Volume* as the outcome. N is very similar across outcomes. t-statistics are presented in parentheses. Standard errors are clustered at the RIC-level. Significance levels indicated as: ** p<10%, ** p<5%, *** 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	22018			21172		

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" spread. While the binary constraint measure above 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 due 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:

$$\begin{aligned}\ln(\text{QuotedSpread}_{it}) = \beta_0 + \beta_1 \ln(\text{Mcap}_{it}) + \beta_2 \ln(\text{Price}_{it}) + \beta_3 \ln(\text{Ntrades}_{it}) + \\ \beta_4 \ln(\text{TradeSize}_{it}) + \beta_5(\text{Volatility}_{it})\end{aligned} \tag{3}$$

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 idiosyncratic volatility. We estimate equation (3) 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 $\text{TickPerSpread}_{\text{it}}$). Then, in the second step, we use the estimated β's from equation (3), 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:

$$Constrained_{it}^{Cont.} = \ln(\text{QuotedSpread}_{it}) - \ln(\widehat{\text{QuotedSpread}}_{it}) \tag{4}$$

In equation (4), 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 $\text{TickPerSpread}_{\text{it}}$ 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 eq. (4), although $\text{TickPerSpread}_{\text{it}} = 2$. 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.[16]

Table 6 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 level of spread constraints. For example, for 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 Order submission strategies

Our main results in Section 3 suggest that the tick size war of June 2009 disproportionately benefited 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 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 economic 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 gauging whether undercutting is more likely than true price competition.

[16]Figure A.15 of the Internet Appendix provides the distribution of $Constrained_{it}^{Cont.}$. In their empirical analysis, Kwan et al. (2015) interact the continuous variable $Constrained_{it}^{Cont.}$ linearly 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 flexibly (and non-linearly) across categories of $Constrained_{it}^{Cont.}$.

Table 6: The market quality effects of the tick size war for constrained and unconstrained stocks — Kwan et al. (2015) constraint measure

The table reports estimates of τ from the following difference-in-differences specification (the same as in Table 4): $y_{it} = \alpha_i + \alpha_t \times \alpha_m + \tau D_{it} + \mathbf{X}_{it} + \omega_{it}$, where α_i, α_m, and α_t represent RIC, market (e.g., Norway), and date-level fixed effects, respectively. $D_{it} = 1$ for all treated group observations on dates $t \geq t^*$, where t^* is the start date of the tick size war for a given trading venue. $\mathbf{X}_{it}$ is a vector of control variables, which includes the natural logarithm of the stock price. The difference-in-differences coefficient τ is estimated separately for the primary exchanges, the MTFs, and the NBBO (Aggregate). The baseline *treated* sample comprises all stocks that could be traded on Chi-X, BATS or Turquoise before June 1, 2009. The *control* sample comprises stocks that could not be traded on Chi-X, BATS, or Turquoise throughout the calendar year 2009. We split the *treated* sample into *Constrained*, *At-Margin* and *Unconstrained* stocks, as defined by the Kwan et al. (2015) procedure described in Section 3.3. All outcome variables except *Volatility*, which is measured in percentage points, have been transformed with the natural logarithm. For ease of exposition, we present N, the number of observations, only for regressions using *Volume* as the outcome. N is very similar across outcomes. t-statistics are presented in parentheses. Standard errors are clustered at the RIC-level. Significance levels indicated as: ** p<10%, ** p<5%, *** 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	21492			21268			20276		

4.1 Undercutting or price competition?

When competing venues offer trading at 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 likelihood 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 achieving that position. Such costs are decreasing in the tick size, with the benefits of prime queue position likely accentuated when tick sizes were previously binding. Accordingly, we may observe single tick undercutting where tick sizes were not particularly binding (as the cost of doing so is 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).[17]

Figure 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). Traders seeking only to undercut the prices at the primary exchange are expected to undercut by one small tick. True price competition would see more activity deeper within the alternative venue's order book.

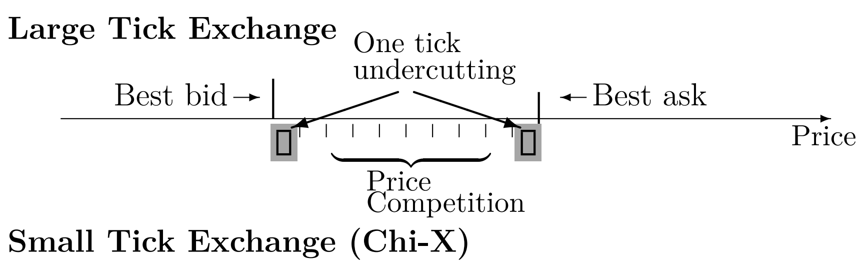

To empirically assess how liquidity providers exploit the access to a new finer pricing grid, we zoom in on the competitive dynamics between the most active alternative trading venue, Chi-X, and the primary exchanges. Figure 4 plots the distribution of the best quotes on Chi-X relative to the primary exchange, measured in alternative-venue tick sizes, one week before and after the start of the tick size war. In Figure 4, "equal" indicates that Chi-X and the primary exchange are quoting identical prices, while +1 (−1) indicates prices on Chi-X

[17]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.

are one new tick better (worse) than the primary exchange.[18]

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 COP and OSE listed stocks. For both COP and OSE, just before the tick size 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 single-tick undercutting was a key driver of the observed increase in Chi-X price improvement during the tick size war.

[18]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.

Figure 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 day 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 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. Copenhagen

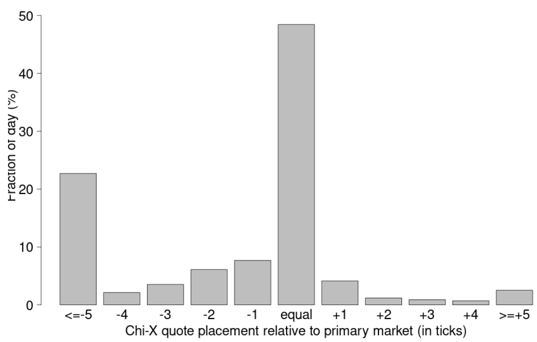

Panel B. Oslo

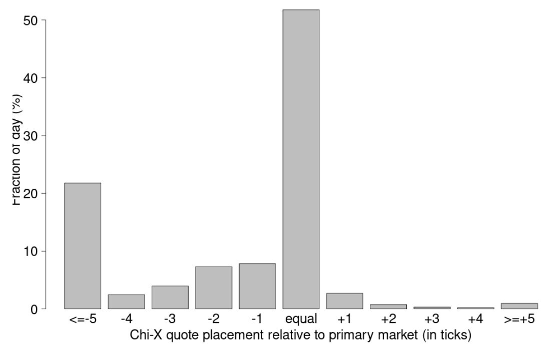

Panel C. Stockholm

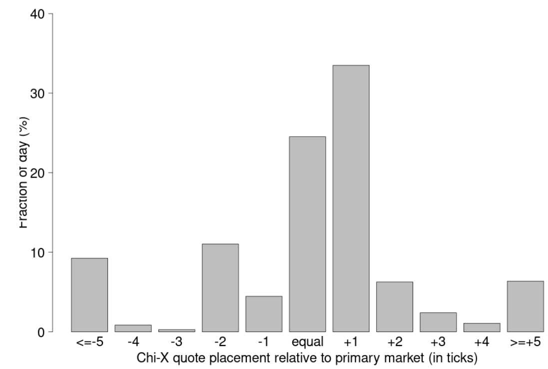

One week post-war

Panel D. Copenhagen

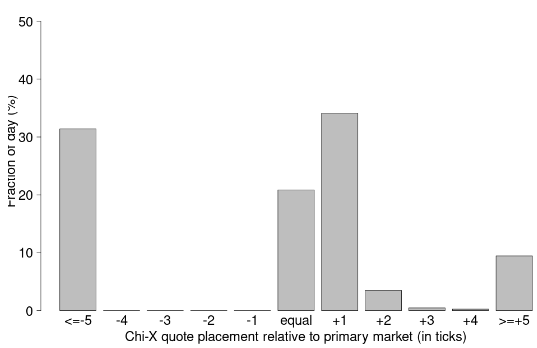

Panel E. Oslo

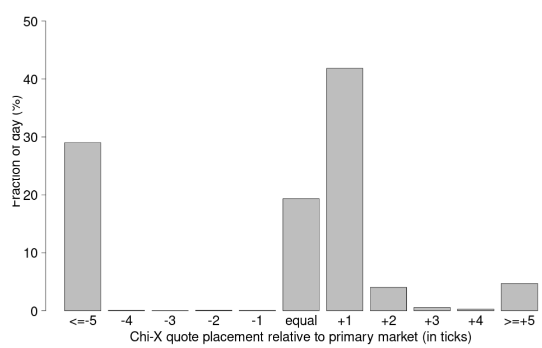

Panel F. Stockholm

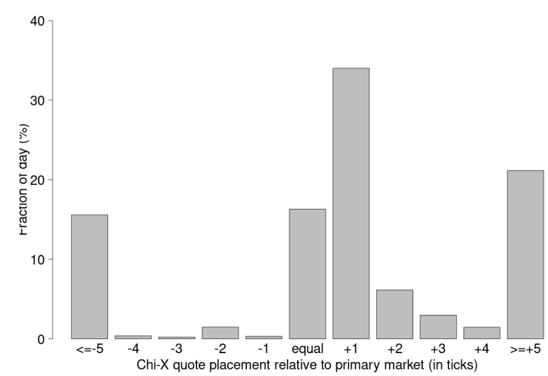

The results for STO stocks, presented in Panels C and F of Figure 4, tell a slightly different story. Unlike for COP and OSE, Chi-X offered smaller tick sizes than STO even before the tick size war. For STO stocks, the tick size war simply widened pre-existing tick size differences. Consequently, just before the tick size war, the Chi-X quotes for STO stocks 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 OSE and COP). After Chi-X further reduced their tick sizes for STO 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 STO stocks. 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 one another increasingly low, resulting in more frequent undercutting, and hence price improvement. Similar undercutting behavior is empirically documented by (Dhyrberg et al., 2019) in the presence of extremely small tick sizes.

4.2 War and Peace: The role of competitive tick sizes

Figure 4 presented 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 fraction 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.[19] 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

[19]Corresponding time series for BATS and Turquoise are presented in an Internet Appendix.

Figure 5: Quote placement at Chi-X relative to primary exchange

The lines 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 4.1. The figure plots weekly averages across all primary exchange stocks with Chi-X trading.

Panel A: Denmark (Copenhagen)

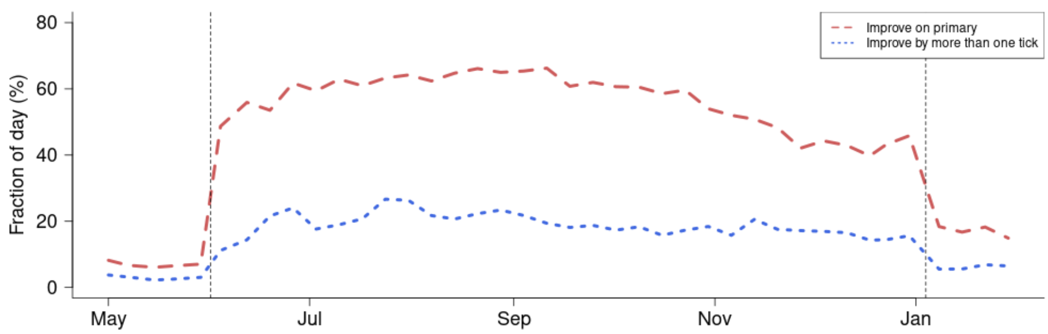

Panel B: Norway (Oslo)

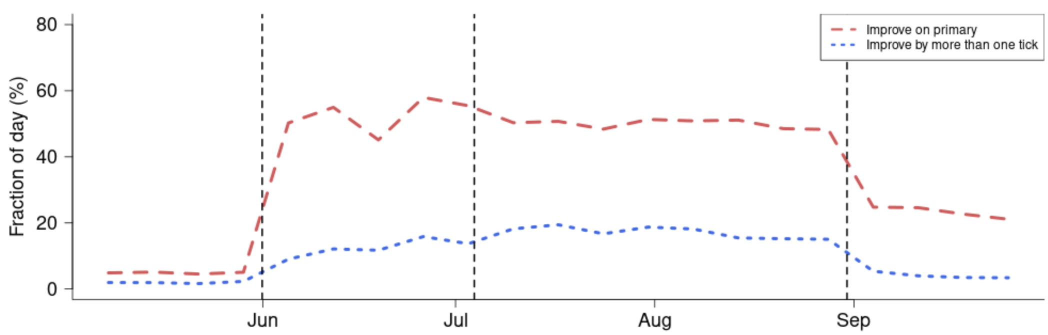

Panel C: Sweden (Stockholm)

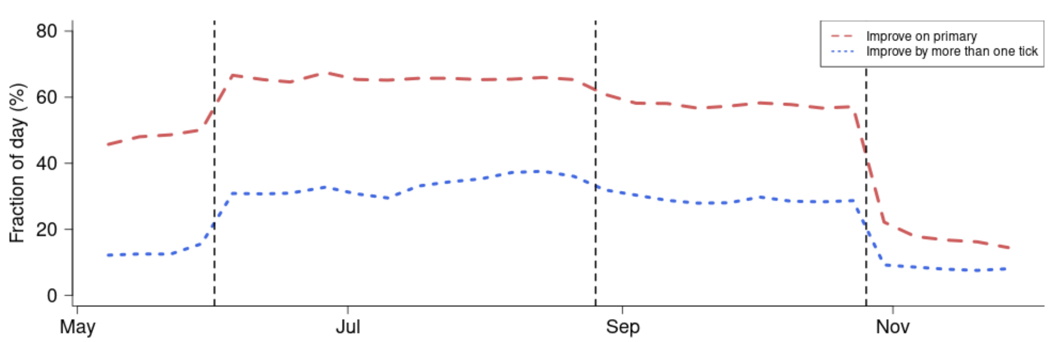

exchange by at least 1 new tick (the red line) to around 60%. The fraction of time for which Chi-X provides price competition of more than 1 new tick (the blue line) jumps by much less on the introduction of the narrower pricing grid, reaching a maximum of only 20-30% of the day. The stark difference in these observed changes are almost immediately reversed to pre-war levels once the tick sizes are harmonized across trading venues.

4.3 Mechanism: What drives price competition?

The time-series evidence in Figure 5 show 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 so may mask any heterogeneity generated by the substantial cross-sectional variation in the level of spread constraint experienced in our sample. The results of our market quality analysis in Section 3 indicate that the competitive benefits of narrower tick sizes were concentrated in constrained securities. We now examine the cross-sectional differences in stock-level 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 narrow price 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 % of time Chi-X improves on the primary exchange by > 1 tick.

(b) The % of time Chi-X improves on the primary exchange by > 3 ticks.

(c) Order book depth at the best Chi-X prices.[20]

[20]Note that the Chi-X depth is measured only at the best prices at Chi-X; it is not the depth measure

The first two variables capture the extent of between-venue price competition — the likelihood of moving beyond 1 (or 3) 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 a number of other factors. Specifically, we estimate the following cross-sectional regression model:

$$y_{it} = \alpha_0 + \beta \text{Constrained}_{i,t} + \gamma \mathbf{X}_{it} + \varepsilon_{it}, \tag{5}$$

where y_{it} is the outcome of interest — i.e., either (a), (b) or (c) as defined above — for stock i on date t and $\mathbf{X}_{it}$ is a vector of control variables, including the stock price and market capitalization. The key explanatory variable is $\text{Constrained}_{i,t}$, which measures whether trading at the primary exchange is constrained by the tick size. We consider three versions of $\text{Constrained}_{i,t}$. The first two are those introduced in Section 3.3 — 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.

Table 7 presents estimates of β from equation (5). The estimates are based on data from June 2009, the first month of the tick size war.[21] The estimates in Table 7 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 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 $\text{Constrained}_{i,t}$. The results are consistent the findings of Foucault and Menkveld (2008), that increased competition — in this case facilitated by a narrower pricing grid — can enhance liquidity provision and simultaneously increase order book depth.

The results presented in Table 7 are also consistent with the market quality improvements documented in Section 3 being concentrated in the constrained securities. When the tick size itself prohibits traders from quoting narrower spreads, market quality is likely to be hindered. The addition of a narrower pricing grid on an alternate exchange facilitates market makers to explicitly compete against each other, reducing the costs of trading. These benefits are

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 low-tick exchange.

[21] Results for the whole tick size war period are available in the Internet Appendix.

Table 7: Are constrained stocks more likely to have Chi-X price competition?

The table reports estimates from regressions of the following form: $y_{it} = \alpha_0 + \beta \text{Constrained}_{i,t} + \gamma \mathbf{X}_{it} + \varepsilon_{it}$, 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)) and (c) The order book depth at the best Chi-X prices (columns (7)-(9)). $\mathbf{X}_{it}$ is a vector of control variables, including the stock price and market capitalization. The table only reports estimates of β. We consider three different measurements of $\text{Constrained}_{i,t}$. *Constrained(binary)* is an indicator variable that equals one if the stock is constrained according to the binary measure introduced in Section 3.3. *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 3.3. *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. The sample uses all stock-days with trading at both Chi-X and the primary exchange. The regressions are estimated as unbalanced panels. t-stats in parenthesis. The standard errors are adjusted for clustering, using the Arellano (1987) adjustment of the White (1980) type of standard errors, as described in Croissant and Millo (2008). Significance levels indicated as: ** p<10%, ** p<5%, *** 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

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 narrower pricing grids may 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 most benefit from a revised pricing grid.

5 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, where three newly entrant 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. Moreover, we find that the competitive tick size reductions significantly improved measures of market quality, such as market-wide spreads and aggregate trading volume. 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 flows 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 U.S. and other jurisdictions (e.g., Comerton-Forde et al. (2019) and Kwan et al. (2015)). However, in cases where a stocks’ trading costs were artificially increased by the tick size — approximately 20% of securities — we find that the introduction of a small-tick venue facilitated more meaningful price improvements, improving on the primary exchange spread by five or more “new” ticks. While there are similarities between “synthetic” mechanisms to avoid a harmonized tick size schedule, such as dark trading or inverted maker-taker, 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 U.S. 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 due to the regulations in both in the U.S. and Europe that enforce a common tick size across trading venues. It is possible that the introduction of dynamic, non-binding tick size schedules which 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 which 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

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.

Manuel Arellano. Computing robust standard errors for within-group estimators. *Oxford Bulletin of Economics and Statistics*, 49(4):431–34, 1987.

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.

Marianne Bertrand, Esther Duflo, and Sendhil Mullainathan. How Much Should We Trust Differences-In-Differences Estimates?*. *The Quarterly Journal of Economics*, 119(1):249–275, 02 2004. doi: 10.1162/003355304772839588.

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.

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 Foucalt. Trading fees and efficiency in limit order markets. *Review of Financial Studies*, 25:3389–3421, 2012.

Jean-Edouard Colliard and Peter Hoffmann. Financial transaction taxes, market composition, and liquidity. *Journal of Finance*, 72(6):2685–2716, 2017. doi: 10.1111/jofi.12510.

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.

Yves Croissant and Giovanni Millo. Panel data econometrics in `R`: the `plm` package, 2008. `R` vignette, available at `CRAN`.

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.

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

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.

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 Meling. Anonymous trading in equities. *Journal of Finance*, 2020. forthcoming.

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 Jhuo Zhong. Relative tick size and the trading environment. *Review of Asset Pricing Studies*, 9(1):47–90, 12 2018. doi: 10.1093/rapstu/ray009.

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

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.

Halbert White. A heteroskedasticity-consistent covariance matrix estimator and a direct test of heteroskedasticity. *Econometrica*, 48:817–38, 1980.

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.