Market Design and Stock Liquidity

Three Essays in Market Microstructure

Tom Grimstvedt Meling

Thesis for the degree of philosophiae doctor (PhD)
at the University of Bergen

2017

Date of defence: September 1, 2017

Acknowledgements

First and foremost, I would like to thank my thesis supervisors, professor Hans K. Hvide and professor Bernt Arne Ødegaard for invaluable guidance and advice throughout the process of producing my dissertation. As my main supervisor, Hans K. Hvide has relentlessly challenged my academic work, pushing me to explore the limits of my intellectual capabilities. I am forever grateful to Hans for the countless times he has read and criticised my work. Bernt Arne Ødegaard joined as my co-supervisor at an early stage, and I am very grateful for his support, skilful guidance, and co-authorship. The collaboration with Bernt Arne has taught me crucial skills from which I will benefit greatly for many years to come.

I wish to thank everyone at the Department of Economics at the University of Bergen for providing a stimulating and enjoyable research environment. I would like to thank the head to the department, professor Jan Erik Askildsen, for facilitating my many research stays abroad. During my time as a PhD student, I spent two months at the Department of Finance at Hong Kong University of Science and Technology (HKUST), ten months at the Haas School of Business, UC Berkeley, and three months at Booth School of Business, University of Chicago. I am very grateful to associate professor Kasper M. Nielsen for hosting my stay at HKUST, and to professor Steven Tadelis for hosting my stay at UC Berkeley. Their hospitality made my research stays immensely enjoyable and educational.

Finally, I would like to thank my family and friends for their support throughout the good and the bad times of producing this thesis. A special thanks to Rannveig for keeping me (somewhat) sane throughout this process.

Abstract

Stock markets allow buyers and sellers to meet and exchange shares in listed companies. The three chapters of my thesis show empirically that subtle changes to the rules that govern the stock trading process can affect the ease with which shares are traded — the liquidity of stock trading. My research into the stock market's microstructure aims to go beyond the existing academic literature by exploring new empirical techniques for causal inference.

In the first chapter of this thesis, I explore a reform at the Oslo Stock Exchange (OSE) to assess the causal effect of trader anonymity on stock liquidity and trading volume. Using a regression discontinuity approach, I find that anonymity leads to a reduction in bid-ask spreads by 40% and an increase in trading volume by more than 50%. The increase in trading volume is mostly accounted for by an increase in trading activity by institutional investors. These results are consistent with theoretical frameworks where informed traders supply and improve liquidity in anonymous markets.

In Chapter 2, Bernt Arne Ødegaard and I show that competitive stock exchanges undercut other exchanges' tick sizes to gain market share, and that this tick size competition increases investors' trading costs. We analyze a recent event where three entrant exchanges, Chi-X, Turquoise, and BATS Europe, reduced the tick size for stocks with an OSE primary listing. We find that the tick size-reducing exchanges captured market shares from the large-tick OSE by influencing the order-routing decisions of high frequency traders. Trading costs at the OSE increased while trading costs in the competing exchanges remained unchanged. Our findings suggest that unregulated stock markets can produce tick sizes that are excessively small.

In the third chapter of this thesis, I assess the causal impact of increasing the tick size on stock liquidity and trading volume in illiquid stocks. Using a regression discontinuity design at the Oslo Stock Exchange, I find that increasing the tick size has no impact on the transaction costs, order book depths, or trading volumes of illiquid stocks. These findings contradict recent theoretical predictions in the market microstructure literature as well as proposals by lawmakers in the United States to increase the tick size for illiquid stocks.

Contents

Introduction

Introduction

Legend has it that Nathan Mayer Rothschild (1777–1836), once the wealthiest man on earth, exploited early knowledge of Wellington's victory at the Battle of Waterloo to speculate on the London Stock Exchange and make a fortune. Some accounts claim that Rothschild's advanced information was delivered by specially-trained carrier pigeons, while other accounts tell of couriers on horseback delivering news of victory faster than Wellington's official dispatch. Whatever the communication device, Rothschild is believed to have had access to valuable information hours or even days before other Londoners.[1]

Centuries later, in 2010, a new $300 million high-speed fiber optic cable was laid between the stock exchanges in New York and the futures exchange in Chicago. The innovation of the new cable was to dig in a nearly straight line — tunneling through the Appalachian mountains instead of circumventing them. The result for users of the new cable was a reduction in the communication time between New York and Chicago by approximately one millisecond. Since its construction in 2010, investments in microwave technology have further reduced communication times between New York and Chicago, leaving the $300 million fiber optic cable obsolete (see Budish et al. 2015).

These two anecdotes illustrate that leveraging speed to exploit information asymmetries has been central to financial markets since their inception. Stock market participants are willing to invest large sums into cutting-edge information technologies, whether carrier pigeons or fiber optic cables, that allow them to access information earlier than others. The reason they do so is that early access to information can be a source of revenue that can more than offset their initial costs of investment. For example, having access to the fastest communications technology between New York and Chicago enables investors to profit from the arrival of new information in Chicago or New York before the information becomes publicly available.

However, the Waterloo trade and the New York – Chicago trade also illustrate an important tension in financial markets: when some market participants have access to more

[1] The details of Nathan Rothschild's alleged speculation on the London Stock Exchange remains a topic of controversy among historians. Ferguson (1998, 2008) argues that Nathan Rothschild indeed received advanced news of Napoleon's defeat at Waterloo but presents evidence that Rothschild's short-term trading gains from this information were smaller than previously thought. Research by the Rothschild family suggests that "[a]lthough it is virtually part of English history that Nathan Mayer Rothschild made 'a million' or 'millions' out of his early information about the Battle of Waterloo, the evidence is slender" (Rothschild Archive 2017).

information than others, the uninformed investors might suffer. This tension between investors arises because any transaction between two parties involves a ‘winner’ and a ‘loser’, atleast in monetary terms. For instance, some unfortunate Londoners had to supply the securities that Nathan M. Rothschild allegedly made a handsome profit from. In other words, the presence of investors with advanced information can make it more costly for other investors to trade in the stock market.

Why should we care that some stock market participants make it more costly for others to trade? The stock market is more than just a zero-sum game where one person’s gain equals the loss of another; it also allows investors with different preferences and investment horizons to meet and exchange shares in listed companies. Some investors may wish to store current wealth for the future, while others may want to reduce risk. These differences between investors opens for mutually beneficial gains from trade. However, the presence of a pigeon-informed Rothschild or a cable-informed New York trader can make it too costly for uninformed investors to participate in the market, which can prevent mutually beneficial stock trades from taking place.

Changing the rules that govern the stock trading process — the stock market’s microstructure — can help reduce the many frictions that prevent stock trades from taking place. In the case of millisecond cable-traders between New York and Chicago, the stock market’s rules can be changed to offer protection to slower investors from the potential exploitation by faster investors. Such protection can be accommodated by market forces, for example by the entry of new stock exchanges that make competition on speed unfeasible, or by market regulations that prohibit competition on speed altogether.[2] Regardless of who accommodates changes to the stock market’s microstructure, it is important that new trading rules accurately address the relevant market imperfections that deter trading, without introducing new and perhaps more serious market imperfections.

The academic community can provide valuable insights into which stock market rules and

[2]Investor demand for protection against ultra-fast ‘high-frequency traders’ has spurred innovation in stock exchanges’ rules and designs. For example, the so-called the Investors Exchange (IEX) was approved by market regulators in the United States as an official exchange in June, 2016. The main innovation of the IEX is to introduce a speed bump — a 350-microsecond slowdown — by coiling 38 miles of cable in order to mitigate high-frequency traders’ speed advantage. The IEX and its founder, Brad Katsuyama, are the focus of Michael Lewis’ best-seller, *Flash Boys*. In his book, Lewis (2014) argues that the ‘market is rigged’ in favor of high-frequency traders, and the IEX speed bumps are portrayed as a necessary market design innovation to level the playing field between investors.

designs are in the public interest, and which are not.[3] Theoretical work strives to provide accurate predictions on the potential impacts of future market design policy decisions. Empirical work attempts to implement procedures that identify causal effects of existing market design features, and tries to extrapolate the empirical findings to future market design policy decisions. In both theoretical and empirical work, the main goal is to identify the economic mechanisms through which stock market rule changes affect market participants' ability to transact.

My thesis contributes to the empirical academic literature on stock market design. In all three chapters, I use empirical methods to isolate the causal impact of three different stock market design choices on measures of stock liquidity — the ease with which investors can exchange shares. In Chapter 1, I find that increasing the extent of trader anonymity facilitates trading between investors and reduces investors' transaction costs. Chapter 2 shows that the tick size — the minimum price increment on a stock exchange — can be reduced by stock exchanges to capture market shares from other exchanges. However, reductions in the tick size worsens stock liquidity in exchanges that keep large tick sizes. In Chapter 3, I show that increasing the tick size has a causal effect on the stock liquidity in already-liquid stocks but has no effect on the stock liquidity in illiquid stocks.

Market regulators who are concerned with promoting liquid stock markets can draw three lessons from my thesis. Chapter 1 suggests to them that allowing stock trading to become less transparent — in essence, exacerbating the potential for asymmetric information — may in fact benefit stock liquidity. This surprising result pertains because certain sophisticated investors, whose trading improves stock liquidity, are more willing to trade when they can do so anonymously. In Chapter 2, Bernt Arne Ødegaard and I argue that imposing regulations on the competition between exchanges can improve trader welfare. Our argument is based on the observation that exchanges' rational decisions to innovate their market designs may reduce the liquidity that is available for stock market participants. Finally, Chapter 3 informs market regulators that increasing the tick size for illiquid stocks may not be the correct market design tool if the goal is to improve stock liquidity.

[3] Researchers have had real impact on a number of market design issues. For example, an influential study by Christie and Schultz (1994) spurred a large-scale investigation into collusive behavior among NASDAQ dealers, which eventually led to considerable market design changes. Regulators, such as the United States Securities and Exchange Commission (SEC), appear to value input from the academic community, and often consult academics before making major market design decisions.

Literature review

Stock markets allow buyers and sellers to meet and exchange shares in listed companies. The three chapters of my thesis illustrate that subtle changes to the rules that govern stock trading can affect the ease with which shares are traded — the liquidity of stock trading. This section provides a non-technical summary of the three chapters of my thesis and connects each of the thesis chapters to the existing academic literature.

Chapter 1: Market transparency

Some stock traders have more precise information than others concerning the fundamental values of listed companies. This superior information can, for instance, be obtained by investing in research into companies' future earnings prospects. Traders who are unable or unwilling to invest in fundamental information may be reluctant to trade shares as they (rightly) suspect they will be on the losing side of the bargain. However, these initially uninformed traders can gradually learn about fundamental values by observing the trading decisions of more informed traders, which can make them more willing to transact.

Stock exchanges can modify the transparency of the stock trading process to influence how quickly and how precisely market participants can observe each other's actions. Transparent stock markets publish abundant information concerning traders' decisions to submit orders to the market (pre-trade transparency) or concerning executed transactions (post-trade transparency).[4] For instance, when trading is pre-trade transparent, market participants may be notified about the prices and quantities specified in all orders submitted to the market. An even more transparent market would reveal the identities of order submitters, either before or after the order is executed.

Stock trading around the world is increasingly conducted in markets that offer little or no transparency. The clearest sign of a move towards low-transparency (or 'opaque') equity trading is the surging popularity of so-called 'dark pools.' These peculiar marketplaces provide traders with complete pre-trade opacity. This means that only the submitting trader has information about an order being placed, and the rest of the market will be informed about

[4] This brief and non-technical survey of the vast equity market transparency literature will largely conflate the issues of post-trade and pre-trade transparency, and instead focus on the main economic mechanisms that relate the extent of market transparency to traders' willingness to transact. For more elaborate discussions of equity market transparency, see for example Foucault et al. (2013), Biais et al. (2005), or Madhavan (2000).

the order only after it executes. Recent figures from the U.S. equity market suggest that dark pools now handle nearly 18% of overall trading activity, up from approximately 4% in 2005 (Securities and Exchange Commission 2015). Similarly, Chapter 1 of this thesis reports that stock exchanges around the world have moved towards more post-trade anonymity, which means they longer publish trader identifiers after completed transactions.

Why are investors attracted to low-transparency marketplaces? For investors whose demands move prices, either because they have access to price-relevant information or because the market's liquidity supply cannot absorb their demand, the answer is simple — transparent markets frustrate their ability to efficiently work large orders because transparent markets expose trader demands and may therefore increase trading costs. Dark markets, in contrast, allow investors to conceal their trading intentions from the market, making it more difficult for other market participants to anticipate and profit from their actions.

Though a lack of transparency should benefit investors who possess valuable information they wish to conceal, other market participants might suffer. The classical theoretical argument in favor of market transparency is that dark markets attract traders with superior information, which exacerbates the information asymmetry in the dark marketplace and reduces investors' willingness to transact (e.g. Huddart et al. 2001). According to this argument, the recent trend towards less transparent stock markets might concern market regulators who wish to maintain a market design that promotes a level playing field among investors.

However, other theoretical work argues that asymmetric information is not the only economic mechanism at play when changing market transparency. For example, Rindi (2008) and Boulatov and George (2013) argue that informed investors frequently trade as liquidity providers, which means that informed investor trading activity can reduce other investors' trading costs. Their theory suggests that reducing stock market transparency can lower the costs associated with trading shares by promoting more trading by informed liquidity providers. Theoretical ambiguity makes the impact of transparency on market outcomes an empirical question.

Chapter 1 of this thesis presents empirical evidence that increasing post-trade trader anonymity — a reduction in market transparency — can reduce investors' trading costs and increase trading volume. An anonymity reform in the period 2008–2010 at the Oslo Stock Exchange (OSE) provides a rare source of exogenous variation. Semi-annually, the

25 most traded stocks at OSE were selected for anonymous trading; all other stocks were not. Comparing just-included and just-excluded stocks in a so-called regression discontinuity design provides causal estimates. I find that anonymity considerably increases stock liquidity and trading volume. For example, relative bid-ask spreads, a standard measure of illiquidity and transaction costs, are 40% lower for anonymous stocks, and trading volume is higher by more than 50%.[5]

The finding that trader anonymity improves stock liquidity and trading volume is consistent with theoretical models that emphasize the benefits of informed liquidity supply in anonymous markets (e.g. Rindi 2008, Boulatov and George 2013). To further explore the empirical support for this class of models, I use detailed transaction-level data on the trading activity of all investors at the OSE. As empirical proxies for 'informed' and 'uninformed' investors, I follow Linnainmaa and Saar (2012) and use institutional and retail investors, respectively. I find that the increase in aggregate trading volume is mostly accounted for by an increase in trading activity by institutional investors while retail investors do not adjust their trading behavior in response to anonymity. I interpret the simultaneous increase in institutional investor trading activity and stock liquidity under anonymity as consistent with informed traders supplying and improving liquidity in anonymous markets.

Chapter 2: Competition between stock exchanges

Competition between stock exchanges is a fairly modern phenomenon. National stock exchanges, traditionally located in nations' capitols, long enjoyed near-monopolist market positions for trading in domestic shares. However, over the last two decades, deregulations and technological advances have reduced the entry barriers for new exchanges to compete with incumbent exchanges.[6] Chapter 2 of this thesis shows that one consequence of competition

[5] The existing empirical evidence is mixed as to whether transparency enhances or degrades stock liquidity. Madhavan et al. (2005) find that stock liquidity declines after increases in pre-trade transparency, while Boehmer et al. (2005) and Hendershott and Jones (2005) find the opposite. Chapter 1 in this thesis summarizes the empirical literature on post-trade transparency, which also has produced ambiguous conclusions.

[6] In the United States, the introduction of Regulation National Market System (or Reg NMS) in 2005 increased the scope for competition among stock exchanges by forcing all stock exchanges to re-route incoming orders to the exchanges currently posting the best prices. Similarly, the 2007 Markets in Financial Instruments Directive (MiFID) regulation in Europe changed the European equity trading industry by abolishing the so-called 'concentration rule' which forced any regulated trade to be executed on the primary market, thereby unleashing competition for European order flow.

is that profit-seeking stock exchanges have incentives to customize their market designs to capture market shares from other exchanges. We also show that such profit-maximizing behavior can increase market participants' trading costs.

Some stylized facts from the trading of shares in Norwegian companies can help understand how fragmented the market for stock trading currently is, and how fast the market became this fragmented. Publicly available data from Fidessa, a data vendor, show that the Oslo Stock Exchange market share of overall trading (including over-the-counter trading) in its most heavily traded stocks declined from 100% in 2007 to nearly 40% by 2016. In 2016, more than twenty regulated markets and unregulated over-the-counter trading venues competed with the Oslo Stock Exchange for investors' order flow. However, the Oslo Stock Exchange remains the largest market for trading in stocks with an Oslo primary listing, followed by BATS over-the-counter, BATS CXE (formerly known as Chi-X), Turquoise, and BATS BXE (formerly known as BATS Europe).

The early theoretical literature considered competition between stock exchanges as an implausible equilibrium because of the perceived network externalities associated with conducting all trading within a single marketplace. For example, in the theoretical models of Pagano (1989) and Chowdhry and Nanda (1991), trading tends to consolidate in a single market in equilibrium because both informed and uninformed traders want to be part of the largest trading crowd. Put bluntly, the mechanism behind this theoretical result is that searching for a suitable counter-party is less costly when there are more potential counterparties in the marketplace.

There are also reasons to think that competition between stock exchanges can make trading easier for market participants. An argument by Harris (1993) is that new exchanges with innovative market designs can be established to cater to the heterogeneous needs of investors, making it easier for these investors to buy or sell shares. For example, the extraordinary success of low-transparency trading venues such as 'dark pools' is perhaps a result of demand from large institutional traders who are concerned with hiding their trading intentions from the public. Moreover, competition may force exchanges to charge smaller trading fees from market participants than a monopolist exchange would (e.g., O'Hara and Ye 2011).

One consequence of being able to trade the same company's shares at multiple exchanges is that high-frequency traders (HFTs) have risen to prominence. While a formal definition of HFT does not exist, HFTs can be thought of as a class of extremely fast computer traders who use complex algorithms to analyze data and execute orders based on their analyses of current market conditions. One of the strategies that HFTs follow is to correct (and profit from) price deviations for the same security at different exchanges. In this endevour, speed is an advantage since it is usually the first mover who gets the best price.[7] In a recent concept release, the Securities and Exchange Commission (2010) noted that HFT activity typically exceeds 50% of total volume in U.S. listed equities and concluded that "[b]y any measure, HFT is a dominant component of the current market structure and likely to affect nearly all aspects of its performance."

Besides arbitraging away between-exchange price discrepancies, HFTs also take on other roles in modern stock markets. In his survey of high-frequency trading, Menkveld (2016) argues that most HFTs operate as market makers, with a business model of providing liquidity to other market participants, being compensated by the bid-ask spread. However, Menkveld (2016) also argues that some HFTs exploit their speed advantage to extract profits from slower traders (what Menkveld (2016) calls the 'high-frequency bandits'). For example, HFTs can react to the arrival of new and valuable information before other traders have time to modify their previous (now mispriced) offers to buy or sell. Other HFT strategies even resemble illegal price manipulation — the so-called 'spoofing' strategy involves filling the order book with buy or sell orders to create the illusion that there is excess supply or demand in the market that should warrant a price movement, only to rapidly modify the orders and profit from investors who erroneously interpret the market's supply and demand as real.

In Chapter 2 of this thesis, Bernt Arne Ødegaard and I explore how stock exchanges customize their market designs to capture market shares from other exchanges, and how such competition affects investors' trading costs. Our empirical setting involves a natural experiment where three stock exchanges, Chi-X, Turquoise, and BATS Europe, in the Summer of 2009 unexpectedly reduced their tick sizes — the smallest price increment on the exchange — for stocks with an Oslo Stock Exchange (OSE) primary listing. The OSE quickly responded by reducing its own tick sizes, before all markets agreed on a common tick size structure.

Consistent with theoretical work by Buti et al. (2015), we find that the tick size-reducing

[7]HFT firms spend vast resources on speed technology to improve their speed advantage. This point was illustrated in the introduction to this thesis with HFTs' willingness to invest $300 million in a new fiber optic cable between the financial markets in Chicago and New York, only to achieve a communication time reduction of approximately one millisecond.

stock exchanges capture market shares from the large-tick OSE. Chapter 2 also shows that between-exchange tick size differences appear to affect the distribution of market shares through their impact on the trading behavior of high-frequency traders. This finding contrasts with Buti et al. (2015), who predict that between-exchange tick size differences affect the distribution of market shares by relaxing constraints to bid-ask spreads in one market. Using a difference-in-differences approach, we also find that trading costs at the OSE increased during the 'tick size war,' while trading costs in the competing exchanges remained largely unchanged.

Chapter 3: Tick sizes

In Chapter 3, I explore whether stock exchanges can promote trading activity and reduce trading costs by changing their tick size — the smallest price increment on the stock exchange. My research in Chapter 3 adds to a recent and fiery debate over the optimal level of the tick size. For instance, in an effort to promote trading in small and illiquid securities, the U.S. Congress has instructed the U.S. Securities and Exchange Commission (SEC) to conduct a large-scale experimental program which will increase the tick size for 1200 small capitalization securities.[8]

Changing the tick size can affect the profitability of different trading strategies, and can therefore affect the dynamics of equity trading. To illustrate this point, consider the theoretical Foucault (1999) model where traders arrive exogenously to the marketplace with private asset valuations. Traders in the Foucault (1999) model can supply liquidity by placing limit orders at some price close to their private asset valuations, or they can demand liquidity by submitting market orders that execute against existing limit orders.

Introducing a positive tick size in this stylized limit order book would make liquidity provision more profitable. This is because the tick size hampers price competition by preventing limit order traders from bidding the asset price up or down to their private valuations. A smaller tick size, in contrast, would intensify competition among liquidity suppliers, since it becomes easier for arriving traders to undercut existing limit orders, which in turn reduces

[8]The 'Tick Size Pilot Program' officially started in October 2016, and will last for two years. The pilot program will increase the tick size for three groups of randomly chosen stocks, and compare changes in stock liquidity for these three groups of stocks with changes in stock liquidity in an unaffected control group of stocks. Each group will comprise approximately 400 securities.

both liquidity providers' profits and liquidity demanders' trading costs.

The real-world effects of changing the tick size are likely to be more complex than portrayed above. For example, when trader arrival is endogenous instead of exogenous, reducing the tick size might induce liquidity providers to exit the market on account of the lower returns to liquidity provision, which negatively affects the market's overall liquidity provision. This has been the pivotal argument in the current U.S. regulatory tick size debate — increasing the tick size for scarcely-traded stocks will increase liquidity providers' profits, promote trader entry, and therefore improve stock liquidity in illiquid stocks. This argument is consistent with recent theoretical work by Werner et al. (2015), who argue that small tick sizes may be optimal for liquid stocks while illiquid stocks may benefit from larger tick sizes.

The established finding in the early empirical literature is that reducing the tick size leads to tighter bid-ask spreads and shallower order books.[9] Motivated by recent theoretical contributions and the current regulatory debate in the U.S., a newer empirical literature explores whether tick size changes affect already-liquid and illiquid stocks differently. Adding to this empirical literature, I assess the causal impact of changing the tick size on stock liquidity and trading volume in both liquid and illiquid stocks. To explore the causal relationship between tick sizes and stock liquidity, I exploit that tick sizes at the Oslo Stock Exchange (OSE) are determined as a function of the stock price — higher priced stocks have larger tick sizes. Comparing stocks that are priced slightly above a tick size price threshold to stocks priced slightly below a price threshold in a regression discontinuity design allows for causal inference. Estimating first the regression discontinuity design for the most liquid stocks at the OSE (stocks in the OBX index), I find that increasing the tick size increases both bid-ask spreads and order book depths. Moreover, I find a weak and potentially time-varying positive impact of tick size increases on trading volume.

To explore whether the effect of tick size changes differs for liquid and illiquid stocks, I apply the regression discontinuity design to a sample which comprises a large number of both liquid and illiquid stocks at the OSE (all non-OBX index stocks). In the period 2008–2011, the 158 stocks in this sample are exposed to more than 2300 exogenous tick size changes, which allows for precise estimation of both average effects and effect heterogeneity. I find that the average causal effect of increasing the tick size for the combined sample of liquid and

[9] For surveys of the academic literature on tick sizes, see for example Holden et al. (2013) and the Securities and Exchange Commission (2012).

illiquid stocks is to widen bid-ask spreads and to increase order book depth. Meanwhile, the average effect is largely accounted for by the most liquid stocks in the sample (top 40% of the liquidity distribution), whose liquidity responds heavily to tick size changes. In contrast, I find no impact of tick size changes on measures of trading costs, order book depth, price volatility, or trading volume for stocks in the bottom 60% of the liquidity distribution.

Bibliography

Biais, B., L. Glosten, and C. Spatt (2005): "Market microstructure: A survey of microfoundations, empirical results, and policy implications," *Journal of Financial Markets*, 8, 217–264.

Boehmer, E., G. Saar, and L. Yu (2005): "Lifting the Veil: An Analysis of Pre-trade Transparency at the NYSE," *The Journal of Finance*, 60, 783–815.

Boulatov, A. and T. J. George (2013): "Hidden and Displayed Liquidity in Securities Markets with Informed Liquidity Providers," *The Review of Financial Studies*, 26, 2096.

Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *The Quarterly Journal of Economics*, 130, 1547.

Buti, S., F. Consonni, B. Rindi, Y. Wen, and I. Werner (2015): "Sub-Penny and Queue-Jumping," Working paper, Ohio State University, Columbus OH.

Chowdhry, B. and V. Nanda (1991): "Multi–market trading and market liquidity," *Review of Financial Studies*, 4, 483–511.

Christie, W. G. and P. H. Schultz (1994): "Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?" *The Journal of Finance*, 49, 1813–1840.

Ferguson, N. (1998): *The House of Rothschild: Money's Prophets, 1798–1848*, Viking Penguin.

——— (2008): *The Ascent of Money*, The Penguin Press.

Foucault, T. (1999): "Order flow composition and trading costs in a dynamic limit order market," *Journal of Financial Markets*, 2, 99–134.

Foucault, T., M. Pagano, and A. Röell (2013): *Market Liquidity. Theory, Evidence and Policy*, Oxford University Press.

Harris, L. (1993): "Consolidation, Fragmentation, Segmentation, and Regulation," *Financial Markets, Institutions & Instruments*, 2, 1–28.

Hendershott, T. and C. M. Jones (2005): "Island Goes Dark: Transparency, Fragmentation, and Regulation," *The Review of Financial Studies*, 18, 743.

Holden, C. W., S. Jacobsen, and A. Subrahmanyam (2013): "The Empirical Analysis of Liquidity," *Foundations and Trends in Finance*, 8, 263–365.

Huddart, S., J. Hughes, and C. Levine (2001): "Public Disclosure and Dissimulation of Insider Trades," *Econometrica*, 69, 665–681.

Lewis, M. (2014): *Flash Boys*, Allen Lane.

Linnainmaa, J. T. and G. Saar (2012): "Lack of Anonymity and the Inference from Order Flow," *The Review of Financial Studies*, 25, 1414.

Madhavan, A. (2000): "Market microstructure: A survey," *Journal of Financial Markets*, 3, 205–258.

Madhavan, A., D. Porter, and D. Weaver (2005): "Should securities markets be transparent?" *Journal of Financial Markets*, 8, 265–287.

Menkveld, A. J. (2016): "The Economics of High-Frequency Trading: Taking Stock," *Annual Review of Financial Economics*, 8, 1 – 24.

O'Hara, M. and M. Ye (2011): "Is market fragmentation harming market quality?" *Journal of Financial Economics*, 100, 459–474.

Pagano, M. (1989): "Trading Volume and Asset Liquidity," *Quarterly Journal of Economics*, 104, 255–274.

Rindi, B. (2008): "Informed Traders as Liquidity Providers: Anonymity, Liquidity and Price Formation," *Review of Finance*, 497–532.

Securities and Exchange Commission (2010): "Concept release on equity market structure," Release No. 34–61358. File No. S7–02-10.

——— (2012): "Report to Congress on Decimalization," SEC Report. Available at http://www.sec.gov/news/studies/2012/decimalization-072012.pdf.

——— (2015): "Shedding Light on Dark Pools, Commissioner Luis A. Aguilar," SEC public statement. Available at https://www.sec.gov/news/statement/shedding-light-on-dark-pools.html.

The Rothschild Archives (2017): "Nathan Mayer Rothschild and 'Waterloo'," Available online at `https://www.rothschildarchive.org/contact/faqs/nathan_mayer_rothschild_and_waterloo`. Accessed on May 9, 2017.

Werner, I., Y. Wen, B. Rindi, F. Consonni, and S. Buti (2015): "Tick Size: Theory and Evidence," Working paper, Ohio State University, Columbus OH.

Chapter 1

Anonymous trading in equities

Anonymous trading in equities

Tom Grimstvedt Meling*

April 3, 2017

Abstract

I explore a reform at the Oslo Stock Exchange to assess the causal effect of trader anonymity on liquidity and trading volume. Using a regression discontinuity approach, I find that anonymity leads to a reduction in bid-ask spreads by 40% and an increase in trading volume by more than 50%. The increase in trading volume is mostly accounted for by an increase in trading activity by institutional investors. These results are consistent with theoretical frameworks where informed traders supply and improve liquidity in anonymous markets.

*University of Bergen, Department of Economics. Tom.Meling@econ.uib.no. This paper has benefited from discussions with Katrine V. Løken, Kasper M. Nielsen, Jialin Yu, Signe Aase Abrahamsen, Carl Nilsen, Xavier Giroud, Hugo Hopenhayn, Jonas Tungodden, Terrance Odean, Christine Parlour, Bernt Arne Ødegaard, Terrance Hendershott, Raffaele Giuliana, Sjur D. Flåm, Øivind Schøyen, Eirik A. Strømland, and especially my thesis adviser Hans K. Hvide. I am also grateful for comments from seminar participants at the University of Bergen, Haas School of Business, Norwegian School of Economics (NHH), and the EEA 2016 Annual Congress. Discussants Oddmund Berg and Konrad Raff have provided invaluable insights. All errors are my own.

1 Introduction

Stock exchanges continually fine-tune their markets to promote liquidity. A much-used strategy in the last decade has been to alter the degree to which traders are anonymous. In this paper, I assess the effect of trader anonymity on stock liquidity and trading volume. An anonymity reform at the Oslo Stock Exchange allows for causal inference. Consistent with theoretical frameworks where informed traders supply and improve liquidity in anonymous markets, I find that trader anonymity increases stock liquidity and trading volume, and that the increase in aggregate trading volume is mostly accounted for by increased activity by institutional investors.

How transparent should trading in equity markets be? Market regulators have long advocated for more transparency. For example, in 2009, the former SEC chairman Schapiro stated that "Transparency is a cornerstone of the U.S. securities market (...) We should never underestimate or take for granted the wide spectrum of benefits that come from transparency" (SEC 2009). Regulators both in the United States and in Europe are currently considering comprehensive market structure changes to increase market transparency.[1]

Market participants, on the other hand, caution that too much transparency frustrates traders' ability to efficiently work large orders because transparent markets expose trader demands and may increase trading costs — thus, harming liquidity.[2] At least partly in response to trader demands, leading stock exchanges, such as Nasdaq, London Stock Exchange, and Deutsche Börse, have recently increased trader anonymity (the Appendix provides an overview of policy changes in this area).[3]

[1] For example, the European MiFID II regulation, due in 2018, is expected to introduce mechanisms that cap the volumes that can be traded in the least transparent venues (European Commission 2014). Similarly, in the United States, the Financial Industry Regulatory Authority (FINRA) recently announced plans to expand its ongoing 'Transparency Initiative' by mandating the public disclosure of block-sized transactions in so-called ATSs, a class of low-transparency trading venues (FINRA 2016).

[2] For example, Øyvind Schanke, head of equity trading at NBIM, the world's largest sovereign wealth fund, recently expressed concerns over transparent markets, to the *Wall Street Journal* (2013): "If we sent our orders into the market, we would have to wait days or weeks for our brokers to execute the trade. Even then, there are risks of information leakage."

[3] In a regulatory appeal to introduce trader anonymity in NASDAQ's SuperMontage system, the exchange stated that: "Nasdaq proposes to add a post-trade anonymity feature to SuperMontage in response to demand from members (...) Anonymity is important to market participants because sometimes the identity of a party can reveal important 'market intelligence' and complicate a member's ability to execute its customer orders" (Federal Register 2003).

Theoretical predictions on the effect of trader anonymity on stock liquidity and trading volume are ambiguous. The literature on liquidity-supplying informed traders (e.g. Boulatov and George 2013, Rindi 2008) posits that informed traders are more willing to supply liquidity when they can do so anonymously. As a consequence, anonymous markets attract liquidity suppliers who improve stock liquidity. In contrast, Huddart et al. (2001) find a negative effect of trader anonymity on liquidity and trading volume because anonymity exacerbates adverse selection which reduces the willingness to transact. Theoretical ambiguity makes the impact of trader anonymity on market outcomes an empirical question.

The purpose of this paper is to empirically assess how trader anonymity affects stock liquidity and trading volume. An anonymity reform in the period $2008 - 2010$ at the Oslo Stock Exchange (OSE) provides a rare source of exogenous variation. Semi-annually, the 25 most traded stocks at OSE were selected for anonymous trading; all others were not. Comparing just-included and just-excluded stocks in a so-called regression discontinuity design provides causal estimates. I find that anonymity significantly increases liquidity and trading volume. For example, relative bid-ask spreads, a standard measure of illiquidity and transaction costs, are 40% lower for anonymous stocks, and trading volume is higher by more than 50%.

Improvements in stock liquidity and trading volume may not be due to trader anonymity but index inclusion effects. Anonymity at OSE was determined by membership in the OBX index, a composition of the most traded shares at OSE. Systematic differences between index and non-index stocks, caused (for example) by index benchmarking strategies, can confound the estimated effects of anonymity. To examine if OBX index stocks are systematically different from non-index stocks, I compare index and non-index stocks in periods before anonymity was introduced. I find no differences between marginal index and non-index stocks in periods without anonymity. Moreover, index funds typically track the broader Oslo benchmark index, in which all sampled stocks are included, and not the OBX by itself. For example, only two index funds track the OBX in the sample period, and their combined net assets amount to 5% of the net assets tracking the benchmark index. Thus, it seems unlikely that index effects are driving the results.

That trader anonymity improves stock liquidity and trading volume is inconsistent with theoretical models that emphasize the adverse selection costs of anonymous markets (e.g. Huddart et al. 2001) but is consistent with models that emphasize the benefits of informed liquidity supply in anonymous markets (e.g. Rindi 2008, Boulatov and George 2013). To further explore the empirical support for the latter class of models, I use detailed transaction-level data on the trading of all investors at the OSE. As empirical proxies for 'informed' and 'uninformed' investors, I follow Linnainmaa and Saar (2012) and use institutional and retail investors, respectively. I find that the increase in aggregate trading volume is mostly accounted for by an increase in trading activity by institutional investors while retail investors do not adjust their trading behavior in response to anonymity. I interpret the simultaneous increase in trading by institutional investors and stock liquidity under anonymity as consistent with informed traders supplying and improving liquidity in anonymous markets.

This paper connects to current debates in the academic literature. First, the existing empirical literature on trader anonymity (e.g. Theissen 2003, Waisburd 2003, Comerton-Forde et al. 2005, Foucault et al. 2007, Thurlin 2009, Comerton-Forde and Tang 2009, Hachmeister and Schiereck 2010, Friederich and Payne 2014, Dennis and Sandås 2015) has produced mixed results. This literature is based on non-exogenous variation where identification is difficult.[4] In contrast, I exploit exogenous variation in trader anonymity for causal inference. My research contributes to this literature with cleanly identified positive effects of trader anonymity on stock liquidity and trading volume.

Second, recent empirical work by Collin-Dufresne and Fos (2015) finds that standard measures of adverse selection and stock liquidity are uninformative about the presence of informed traders. They argue that the classical inverse relation between informed trading and stock liquidity breaks down, among other reasons, because informed traders supply liquidity. My results complement their findings by showing a positive relationship between informed trading activity and stock liquidity in anonymous markets.

Moreover, my research may provide guidance to policy makers in the United States and Europe who are currently considering market structure changes to increase equity market transparency (see footnote 1). The results in this paper suggest to them that increasing equity market transparency may worsen overall market quality by discouraging informed

[4] The existing literature on trader anonymity is based mostly on between-market comparisons and before-and-after variation in anonymity, which does not allow for a separation of the effect of trader anonymity from confounding factors. Recent studies use difference-in-differences strategies with different markets as control groups to improve identification (Dennis and Sandås 2015, Friederich and Payne 2014). This variation is unlikely to be exogenous, as the choice to implement anonymity for a given market is likely to be endogenous to its future market quality trend.

investors from participating and providing liquidity to the market.

The paper is organized as follows: Section 2 presents the anonymity reform at the Oslo Stock Exchange; Section 3 describes the data; Section 4 describes the empirical design; Section 5 presents the main results; Section 6 investigates the validity and robustness of the empirical design; Section 7 explores the mechanisms driving the main results; and Section 8 concludes.

2 The reform

This section begins with a brief presentation of the Oslo Stock Exchange before providing details on a trader anonymity reform from $2008 - 2010$ at the Oslo Stock Exchange.

2.1 The Oslo Stock Exchange

The Oslo Stock Exchange is a medium-sized stock exchange by European standards, currently ranking among the 35 largest equity markets in the world by market capitalization (World Federation of Stock Exchanges 2016). At the end of 2010, the total market capitalization of OSE was about 1.7 trillion NOK (1USD $\approx$ 8NOK), spread out over 220 companies.[5] Turnover velocity in the period $2008 - 2010$ ranged between 124.9% and 156.8%. Faced with competition from alternative trading venues, OSE market shares for trading in OSE listed stocks declined from 100% in 2008 to approximately 90% in 2010. By 2015, this figure is close to 60%.[6]

The OSE operates a fully electronic centralized limit order book and has done so since 1999. The OSE order book allows conventional limit, market, and iceberg orders, along with various other order types. As is common in electronic order-driven markets, order placements follow price-time priority: orders are first sorted by their price and then, in case of equality,

[5] These figures are extracted from Oslo Stock Exchange annual statistics, publically available on the OSE web site. Statistics on the trading of OSE listed stocks on alternative trading venues are based on publically available data collected from Fidessa, a data vendor.

[6] The Oslo Stock Exchange has, in recent years, been the testing-ground of several empirical studies. For example, Ahern and Dittmar (2012) use a Norwegian gender quota reform to investigate the impact of female board representation on firm valuations at the Oslo Stock Exchange. Døskeland and Hvide (2011) leverage high-quality administrative data to investigate the trading performance of individual investors in professionally close stocks, while Næs et al. (2011) explore the connection between stock market liquidity and the business cycle.

by the time of their arrival. The trading day at the OSE consists of three sessions: an opening call period, a continuous trading period, and a closing call period. In late 2012, the continuous trading session was shortened from 09:00 – 17:20 to 09:00 – 16:20. Call auctions may be initiated during continuous trading if triggered by price monitoring or to restart trading after a trading halt.

2.2 Trader anonymity at the Oslo Stock Exchange

The Oslo Stock Exchange (OSE) introduced post-trade anonymous trading on June 2, 2008.[7] Anonymity was introduced to the 25 most traded stocks at the OSE — the constituents of the OBX list. The OBX list is aimed to be a highly liquid composition of shares that reflects the Oslo Stock Exchange investment universe. The stock composition of the OBX list has been revised twice a year (end of June and December) since 1987. After June 2, 2008, all OBX stocks were traded anonymously, and all other stocks at OSE were traded non-anonymously. Stocks entering the OBX after this date received anonymity, and stocks leaving the OBX lost anonymity. See Figure 1 for a time-line.

Stocks are selected for the OBX list based on cumulative trading volume in the six months leading up to a new OBX composition. Table 1 illustrates the selection process. On all list revision dates in Table 1, the 25 stocks with the highest currency trading volume accumulated over the previous six months are chosen from the broader Oslo benchmark index (OSEBX) to comprise the OBX for the subsequent six months.[8] If, for example, two stocks X and Y have accumulated trading volumes of 10 billion NOK and 10.1 billion NOK, respectively, then stock Y is ranked above stock X and is more likely to become an OBX stock. If both stocks rank among the 25 most traded, they will both become OBX listed stocks. If, however, stock Y is ranked 25, and stock X is ranked 26, the former will be an OBX stock, and the latter will not.

[7] The Oslo Stock Exchange often consults members before making major changes to the market model. Members were consulted on whether to introduce trader anonymity or not in a letter dated April 2007. The consultation response was only slightly in favor of implementing anonymity, which may explain why anonymity was implemented only for a small group of the stocks. The decision to implement anonymous trading was first announced February 19, 2008.

[8] The OSEBX is the benchmark index at the OSE. The OSEBX index is an investible index which comprises the most traded shares of the Oslo Stock Exchange. It is revised semi-annually on a free-float adjusted basis. Revisions of the OSEBX index take place on 1 December and 1 June. The OSEBX index typically holds between 60 and 80 stocks, from which the 25 OBX list stocks are chosen.

OSE can supersede the volume-based assignment procedure if “special circumstances so indicate.” When the OSE chooses to do so, there is a disparity between the predicted assignment and actual assignment, which needs to be accounted for in the empirical application. A stock may, for example, be exempt from the semi-annual volume-ranking if trading frequency is too low, turnover is too volatile, or the stock is intended for delisting from the exchange. If the OSE chooses to override the main assignment rule, it fully excludes the stock from the ranking process due to non-eligibility.

The trader anonymity introduced by the OSE significantly reduced the amount of information disclosed from the trading process. The top panel of Table 2 illustrates the information available to market participants when trading is non-anonymous. All market participants observe the identities of buyers and sellers (at the brokerage firm level) instantaneously after transactions, in addition to prices and volumes. In contrast, when trading is anonymous, this information is no longer available (bottom panel of Table 2).[9] Market participants observe that transactions have been executed, with corresponding prices and volumes, but do not observe the identities of buyers or sellers.

Transparency was restored for all stocks after two years. On April 12, 2010, the OSE adopted a new trading platform and, at the same time, reversed the trader anonymity rule. Therefore, trading in all stocks at the OSE is currently fully transparent and non-anonymous.

3 Data

Data are collected from several sources. I collect daily frequency data on all common stock at the Oslo Stock Exchange from *Børsprosjektet* at the Norwegian School of Economics (similar to CRSP). The data covers the period December 2001 - December 2010. This dataset holds information on opening and closing prices, daily price dispersion (highest and lowest prices), measures of trading volume (in currency and in shares), end-of-day bids and asks, and OBX

[9]Identities were available in real time bilaterally to the parties of the trade, and to all market participants after the close of each trading day (daily batch updates at 18:00). The OSE introduced a central clearing party (CCP) in June 2010 after both the introduction and reversal of trader anonymity. This means that, in order to facilitate clearing and settlement, the identities in each specific transaction had to be disclosed to the specific counterparty of the transaction, even with anonymous trading. The anonymity reform implies a move from multilateral to bilateral exposure of identities.

and OSEBX index constituency indicators.[10] I supplement this data with the daily number of transactions, obtained from the OSE. I use these data to assess the impact of trader anonymity on market quality (Section 5).

From *Børsprosjektet*, I also collect yearly frequency data on a variety of firm characteristics and accounting measures. This dataset contains information on firms' total equity, total assets, market capitalization, price-to-book ratio, operating profits, operating income, and cash holdings. Firm characteristics are collected on the last trading day of each calendar year.[11] I use these data to assess whether just-included and just-excluded OBX stocks are comparable in their observable characteristics.

In the analysis on heterogeneity in trader response to anonymity (Section 7), I use proprietary transaction-level data obtained from the OSE. The data contains time-stamped (to the nearest second) information on all transactions in all common stock at the OSE. Each entry in the dataset is a trade and gives the identity of buyers and sellers as well as volumes, prices, and stock identifiers. Trader identifiers were not available to market participants in this period, but the OSE kept record for market surveillance purposes. Buyer and seller identities are at the brokerage level and do not identify underlying accounts.

3.1 Sample selection

In the main analysis (Section 5), I investigate the effects of trader anonymity at the Oslo Stock Exchange and restrict the sample period to June 2008 − April 2010. In falsification tests (Section 6), I employ the full sample period, from 2002 − 2010, to analyze revisions of the OBX list both before and after trader anonymity was introduced. In both analyses, I restrict the sample to the 70 trading days following each OBX revision date. Relevant OBX revision dates are found in Table 1. These 70-day trading windows are defined as events and identified by subscript e. This restriction is imposed to ensure that each event is of equal duration, as transparency was restored April 12, 2010, between OBX revision dates.

[10] Due to minor errors in the OBX constituent data from *Børsprosjektet*, data on OBX list constituency have been corrected using hand-collected data from electronic archives at the OSE. Historical data on tick sizes have been compiled from the same source.

[11] While some of the firm characteristics, such as market capitalization and price-to-book, may be defined on a higher frequency, for simplicity, I define all firm characteristics on the same, yearly frequency. In order to assign firm characteristics and accounting variables to firms that are delisted from the OSE during the calendar year, I collect (from *Børsprosjektet*) a weekly frequency dataset containing the same set of firm characteristics and assign characteristics to firms on the final observation date before delisting.

The transaction-level data used in Section 7 covers four weeks of trading following each of the four OBX revisions in $2008 - 2010$. For balance, I restrict the sample to the 16 trading days following each revision (analogously defined as an event e). As is customary with transaction-level datasets, I keep only automatically matched on-order book trades that are executed during normal exchange opening hours. When, in Section 7, I compute the number of trades and trading volume for different investors, I only consider buy transactions to avoid double-counting transactions.

In sections $5 - 7$, I collapse the data at the event-level. Variables are first defined on a daily frequency, then averaged within each event e. For example, the log of number of trades is defined daily as $lntrades_{it}$ for stock i on date t and averaged into a single observation $lntrad_{ie}$ for event e.[12] I do this to ease the intuition of the regression discontinuity design, which is often associated with cross-sectional data, applied throughout the analysis.

Throughout the analysis, I only keep stocks listed on the benchmark index at the OSE (the OSEBX index). Only OSEBX stocks are eligible for the semi-annual volume-ranking that determines OBX list constituency and, consequently, anonymous trading. The OSEBX index usually holds $60 - 80$ stocks, from which the 25 OBX list stocks are chosen.

3.2 Summary statistics

Table 3 summarizes stock characteristics in the full sample period $2002 - 2010$. Two features of the data stand out. First, OBX listed shares are (on average) vastly different from other shares listed at the Oslo Stock Exchange across all observable characteristics. For example, OBX shares are significantly more valuable, more frequently traded, and have lower transaction costs than non-OBX shares. This is the natural consequence of the volume-based OBX list selection mechanism.

Second, the sampled stocks are mostly small- or medium-capitalization firms, by international standards. For example, the average firm market capitalization is 18.6 billion NOK (1 USD $\approx$ 8 NOK), which is comparable to large S&P600 (small-cap) stocks or small S&P400 (mid-cap) stocks. The stocks are, however, actively traded. The average share volume is 1.6

[12] The stock panel is not balanced because some stocks are delisted from the Oslo Stock Exchange before the 70 day event window is over. For these stocks, outcomes are computed using the number of trading days available. Applying the regression discontinuity design to the full panel of daily observations, instead of on event-level averages, produces almost identical results.

million shares, with a standard deviation of 6.6 million shares. The average stock-day has 451 transactions and a monetary trading volume of 81 million NOK. The average trade size is 4327 shares, and the average trade value is greater than 150 000 NOK.

4 Methodology

I wish to estimate the causal impact of trader anonymity at the Oslo Stock Exchange on stock outcomes. The 'naïve' regression compares outcomes y_{ie} (e.g. stock liquidity) for anonymously traded stocks and non-anonymously traded stocks:

$$y_{ie} = \alpha + \gamma D_{ie} + u_{ie},$$

where D_{ie} is an indicator for anonymous trading in stock i during event e. The effect of interest is captured by the coefficient γ, while the error term u_{ie} represents all other determinants of the outcome. While straightforward to derive, the coefficient γ is unlikely to represent the causal impact of trader anonymity on outcomes y_{ie}. The reason for this is that only the most traded stocks at the Oslo Stock Exchange are traded anonymously such that D_{ie} is likely to be correlated with omitted variables that are themselves correlated with y_{ie} — causing a biased estimate of γ.

The rank-based anonymity assignment mechanism at the Oslo Stock Exchange provides a source of exogenous variation that can be used to overcome this endogeneity problem. The 25 most traded stocks at the OSE are semi-annually assigned to anonymity, while stocks ranked 26 and below are not. Lee (2008) demonstrates that comparing just-included and just-excluded stocks provides quasi-random variation in anonymity since, for narrowly decided races, the outcomes are unlikely to be correlated with other characteristics as long as there is some unpredictable component of the ultimate rank outcome.

The regression discontinuity (RD) design exploits this quasi-random variation (see Lee and Lemieux 2010 for a review). The RD relates discontinuities in outcomes at some treatment threshold to discontinuities in the probability of treatment at the same point. In the case of trader anonymity at the Oslo Stock Exchange, the RD approach implies comparing stocks that are ranked (by previous six month trading volume) marginally inside the top 25 to those ranked marginally outside the top 25.

The first step in the RD design is to define the mechanism that determines eligibility to anonymous trading. I generate a variable r_{ie} that ranks all stocks (1 highest, n lowest) based on the total trading volume in the six-month turnover period leading up to event e. This variable is updated on each OBX list announcement date in the period $2002 - 2010$ (see Table 1). Stocks with a ranking, r_{ie}, at or below the threshold, 25, are predicted for anonymous trading by the main assignment rule:

$$T_{ie} = \mathbf{1}\left[r_{ie} \leq 25\right],$$

where T_{ie} is an indicator variable equal to one for stock i predicted to be traded anonymously after revision e. I normalize the ranking variable by subtracting r_{ie} from 25. The assignment rule becomes:

$$T_{ie} = \mathbf{1}\left[r_{ie} \geq 0\right]. \tag{1}$$

The second step is to identify the relationship between the predicted treatment T_{ie} and actual treatment D_{ie}. In my setting, there is a disparity between T_{ie} and D_{ie} due to non-compliance to the assignment rule 1. I account for this disparity by using a two-stage least-squares procedure (2SLS). Intuitively, the 2SLS approach identifies a discontinuity in the probability of treatment, exactly at $r_{ie} = 0$, and uses this discontinuity to scale any discontinuities in y_{ie} at the same point. The first stage regression can be stated as:

$$D_{ie} = \alpha_0 + \phi r_{ie} + \psi T_{ie} + \omega T_{ie} \times r_{ie} + \varpi_{ie} \tag{2}$$

Since r_{ie} is centered on zero, its inclusion as a regressor in equation 2 ensures that all identification is centered on $r_{ie} = 0$. Notice that if $\psi = 1$, then T_{ie} perfectly predicts D_{ie}, and the probability of treatment jumps from zero to one at $r_{ie} = 0$. Since there is non-perfect compliance to the assignment rule, the coefficient ψ will be less than one.[13] It is the magnitude of ψ that distinguishes this 'fuzzy' RD design from a 'sharp' RD design.

Finally, the second stage regression relates outcomes y_{ie} to treatment status D_{ie} and the ranking variable r_{ie}:

$$y_{ie} = \alpha_1 + \nu r_{ie} + \tau D_{ie} + \beta D_{ie} \times r_{ie} + \varepsilon_{ie}. \tag{3}$$

[13]Estimates from the first-stage relationship in equation 2 are discussed in detail in appendix A.2.

The coefficient τ identifies a discontinuous change in y_{ie} exactly at $r_{ie} = 0$, properly scaled by the first stage relationship. This coefficient can be interpreted as the causal effect of trader anonymity on y_{ie}, under the identifying assumption that stocks are comparable on both their observable and unobservable stock characteristics at $r_{ie} = 0$.[14]

While it is impossible to assess whether stocks close to $r_{ie} = 0$ are similar in their unobservable characteristics, it is straightforward to assess whether or not they are similar in their observable characteristics. In Figure 3 I plot observable stock characteristics over r_{ie}, for all realizations of r_{ie} in the period $2002 - 2010$. The figure shows that all stock characteristics evolve smoothly across $r_{ie} = 0$. This implies that observations close to $r_{ie} = 0$ are, at the very least, comparable in their observable characteristics.

Moreover, the data allow for a powerful falsification test of the RD design. Out of all the realizations of r_{ie} in the period $2002 - 2010$, only the realizations of r_{ie} in the period $2008 - 2010$ actually assigned trader anonymity to OBX listed stocks. This enables me to estimate the coefficient τ both before and after trader anonymity was implemented. Doing so, I document non-zero estimates of τ exclusively in periods with trader anonymity. This addresses a justified concern of simultaneous shocks to y_{ie} at $r_{ie} = 0$. Particularly, if OBX constituency by itself is correlated with outcomes, then estimates of τ are biased. My falsification test, however, suggests that there is no OBX constituency effect in periods without trader anonymity.

The unbiased estimation of τ requires a strong assumption about the functional form of the relationship between r_{ie} and y_{ie}. This assumption is required because, in order to estimate the effects that occur close to $r_{ie} = 0$, it is necessary to use data away from this point as well (Lee and Lemieux 2010). The RD literature has proposed two main approaches to estimating equation 3 when the functional form of r_{ie} is unknown. The first approach, which is widely preferred, is to restrict the sample size on either side of $r_{ie} = 0$ and estimate equation 3 non-parametrically with so-called local linear regressions. If there is a concern that the regression function is not linear over the entire range of r_{ie}, restricting the estimation range to values closer to the cutoff point $r_{ie} = 0$ is likely to reduce biases in the RD estimates (Hahn et al. 2001, Lee and Lemieux 2010). In contrast, the second approach uses all the available data and allows for a flexible relationship between y_{ie} and r_{ie} by expanding equation 3 with polynomials in r_{ie}.

[14]Figure 2 provides a graphical illustration of the 'fuzzy' regression discontinuity design.

I estimate equation 3 non-parametrically with local linear regressions. This implies estimation within so-called bandwidths. In my setting, the bandwidth is the number of stocks included on either side of the treatment cutoff $r_{ie} = 0$. For example, if the bandwidth is $h = 15$, this implies estimating equation 3 for a sample of stocks ranked $r_{ie} \in [-15, 14]$. For transparency and robustness, I present estimates from a wide range of bandwidths.

Tick sizes, the minimum pricing increment, are determined differently for anonymous and non-anonymous stocks and have been found to affect stock characteristics — in particular, stock liquidity (e.g. Bessembinder 2003, Buti et al. 2015). For this reason, I include $ticksize_{ie}$ as a control variable in all specifications.[15]

I follow Card and Lee (2008) and cluster standard errors at the level of r_{ie}.

5 Main results

In this section, I estimate the impact of trader anonymity on stock liquidity and trading volume in the period $2008 - 2010$. The theoretical literature on liquidity-supplying informed traders (e.g. Boulatov and George 2013, Rindi 2008) posits that informed traders are more willing to supply liquidity when they can do so anonymously. Consequently, anonymous markets attract liquidity suppliers who improve stock liquidity. In contrast, Huddart et al. (2001) posit a negative effect of trader anonymity on liquidity and trading volume because anonymity exacerbates information asymmetries which reduce the willingness to transact. Estimates of the empirical effect of trader anonymity at the Oslo Stock Exchange are presented in Table 4.

5.1 Results

I first investigate how trader anonymity affects stock liquidity. I measure stock liquidity with the natural logarithm of relative bid-ask spreads (end-of-day quotes divided by the quote midpoint). Wider bid-ask spreads imply lower stock liquidity and higher transaction costs.[16]

[15]Tick sizes at the OSE are determined as step functions of prices such that higher prices give higher tick sizes. The price cutoffs that determine tick sizes are different for OBX and non-OBX stocks.

[16]The end-of-day relative spread is a crude measure of stock liquidity. The effects documented with this liquidity measure, however, also hold for high-frequency measures of liquidity. For example, in unreported regressions, I evaluate the impact of trader anonymity on common measures of liquidity, such as effective and realized spreads, and document similar effects. I only have access to high-frequency measures of liquidity in

Trader anonymity causes a marked reduction in bid-ask spreads. The estimated effect ranges from -0.86 log points (-58%) to -0.56 log points (-43%), depending on the bandwidth choice. All coefficients are highly significant both statistically and economically. Estimates stabilize at lower levels for larger bandwidths (see also Figure 4 for a richer set of bandwidth specifications).

A second question is whether trader anonymity has any effect on trading behavior. If traders engage in the same transactions irrespective of the anonymity of the trading process, then a reduction in bid-ask spreads simply redistributes revenue from liquidity suppliers to liquidity demanders and has no impact on aggregate welfare. To detect any changes in trading behavior, I estimate the impact of trader anonymity on trading volume, measured both by the number of transactions and currency volume traded. The estimated effect of trader anonymity on log(number of trades) ranges between 0.99 log points ($h = 10$) and 0.51 log points ($h = 20$) with t-statistics between 2.63 ($h = 15$) and 3.35 ($h = 20$). Similar effects are found for the log of value traded. All estimates are statistically significant and imply a tremendous willingness to trade anonymous stocks, relative to transparent stocks.

As an additional test of the impact of trader anonymity on the quality of equity trading, I investigate how anonymity affects the efficiency of prices, proxied by close-to-close returns volatility.[17] Greater volatility is viewed as a trading friction such that the lower the volatility, the more efficient the market. Table 4 shows that anonymous trading has no impact on this measure of price efficiency.

In the appendix of this article, I propose several extensions to the baseline RD model and show that the results in Table 4 are robust. First, I show that the results are not driven by a functional form assumption on the relationship between outcomes y_{ie} and the ranking variable r_{ie}. The results hold for a wide range of polynomials in r_{ie}. Second, I follow Cellini et al. (2010) and Cuñat et al. (2012) and expand the static RD design into a dynamic RD design. The RD design in equation 3 is static in the sense that it does not take into account that anonymous trading in one period potentially affects the probability of receiving

the 'treatment' period $2008-2010$ and not in the 'placebo' period $(2002-2007)$. For comparability between sample periods, I use the end-of-day bid-ask spread throughout the analysis.

[17]My approach is to compute returns volatility for each stock as the sample variance of the close-to-close returns process within each event e. In contrast, much of the existing empirical microstructure literature focuses instead on high-frequency within-day measures of volatility. In unreported regressions, I use a within-day measure of price dispersion — the daily high price divided by daily low price — and the inference remains identical.

anonymous trading in subsequent periods. Such dynamics can arise because 1) anonymous trading is assigned based on trading volume and 2) anonymous trading increases trading volume. In the appendix, I show that the results are not driven by dynamics.

5.2 Summarizing the results

The results in this section suggest that trader anonymity improves stock liquidity. Estimates from a regression discontinuity design show that trader anonymity causes a reduction in bid-ask spreads of more than 40% and an increase in trading activity (trades and trading volume) of more than 50%. These benign effects of trader anonymity on the quality of trading cannot be reconciled with theoretical models that emphasize the adverse selection costs of anonymous markets. Instead, the results appear consistent with theoretical models that emphasize the role of informed liquidity suppliers in anonymous markets.

6 Identification concerns

The previous section established a positive relationship between trader anonymity at the Oslo Stock Exchange and measures of stock liquidity and trading volume. In this section, I discuss whether these relationships can be interpreted as causal. Supportive of a causal interpretation, I find non-zero regression discontinuity estimates exclusively in periods with trader anonymity and not in 'placebo' periods without trader anonymity. Moreover, I show that the effects documented in Section 5, do not seem to be driven by time-varying confounders.

For expositional purposes, I henceforth report estimates only from bandwidth specification $h = 15$.

6.1 Index inclusion effects

The main identification concern in my setting is index inclusion effects. Trader anonymity at the Oslo Stock Exchange was determined by membership in the OBX index. Consequently, all empirical specifications so far have represented joint tests of the effect of trader anonymity and OBX index constituency. If index constituency by itself is correlated with outcomes —

for example due to index benchmarking strategies — my estimates of the effect of trader anonymity may be confounded.[18]

To examine if OBX index stocks are systematically different from non-index stocks, I compare index and non-index stocks in periods before trader anonymity was introduced. Particularly, I exploit that the full sample covers all OBX index revisions in the period $2002-2010$ and that only the index revisions in the sub-period $2008-2010$ assigned trader anonymity to OBX stocks. In column two of Table 5, I apply the baseline regression discontinuity design to data from the 'placebo' period $2002-2007$.[19] I find no differences between marginal index and non-index stocks in periods without trader anonymity.

To formally quantify the difference in regression discontinuity estimates between the trader anonymity period and the 'placebo' period, and to improve statistical precision, I pool all the data and estimate the following difference-in-differences model:

$$y_{ie} = a + \underbrace{\nu r_{ie} + \tau D_{ie} + \gamma r_{ie} \times D_{ie} + \delta ticksize_{ie}}_{Baseline\,model} + \underbrace{\theta ANON_{ie} + \tau^{Diff} D_{ie} \times ANON_{ie}}_{Added\,terms} + \varepsilon_{ie}, \tag{4}$$

where r_{ie} and D_{ie} are defined as earlier. $ANON_{ie} = 1$ for the anonymity period $(2008-2010)$ and 0 for the 'placebo' period $(2002-2007)$, and controls for level differences in y_{ie} between the two periods. The coefficient τ now represents the regression discontinuity estimate in the 'placebo' period. Consequently, the coefficient τ^{Diff} gives the added effect of OBX index constituency in the period $2008-2010$ relative to the period $2002-2007$. Estimates of τ and τ^{Diff} are presented in column three of Table 5. Estimates of the 'placebo'

[18]The literature points to several reasons as to why index stocks might be different from non-index stocks. For example, Boone and White (2015) find that just-included Russell 2000 index stocks have higher liquidity and trading activity and lower information asymmetries than just-excluded stocks. They argue that this is due to greater institutional ownership driven by indexing and benchmarking strategies. A substantial literature shows how index inclusion leads to pricing effects due to excess demand from passive funds tracking the index (see Shleifer, 1986, Harris and Gurel, 1986, Chang et al. 2014). Moreover, limited investor attention could cause salience such that indexed stocks are more heavily traded (see Barber and Odean, 2008, Hirshleifer et al., 2009).

[19]Ideally, I would apply the regression discontinuity design to placebo periods both before anonymity was introduced and after transparency was restored. However, shortly after the OSE restored transparency for all stocks, the exchange introduced new trading rules differentiated between OBX and non-OBX stocks. For example, a central clearing party (CCP) was introduced for OBX stocks only, and new rules for hidden liquidity, differentiated between OBX and non-OBX stocks, were implemented. For this reason, the placebo sample only covers OBX index revisions in the period $2002-2007$.

regression discontinuity effect (τ) remain statistically indistinguishable from zero for all outcome variables, suggesting that marginal OBX and non-OBX stocks are comparable in the absence of trader anonymity. Moreover, the table shows that coefficient estimates of τ^{Diff} are quantitatively similar to estimates from the baseline specification (column 1), although now estimated with significantly more precision.

While the difference-in-differences model efficiently addresses the concern of a fixed index confounder, it does not address potentially time-varying index confounders. In particular, index benchmarking strategies have grown in popularity over the last decade (e.g. Chang et al. 2014).[20] The impact of such a trend might reveal itself through the absence of an index effect in early periods and the existence of one in later periods, which is consistent with the results in Table 5. In an attempt to address such a confounding trend, I conjecture that an increase in index benchmarking only affects the stocks that actually move in or out of the OBX index and not the stocks that remain inside or outside of the index. Therefore, I add to specification 4 separate indicator variables for stocks that move in or out of the OBX, following a revision. This approach allows me to separate any excess effect for moving stocks, from the direct effect of trader anonymity. Column five in Table 5 shows that coefficient estimates and statistical significance are unaffected by the inclusion of mover dummies.

Three institutional details may explain why index constituency seems to have little impact on the results in this paper. First, the bulk of index funds track the broader Oslo benchmark index (OSEBX), in which all sampled stocks are included, and not the OBX by itself. For example, only two index funds track the OBX in the sample period, and their combined net assets amount to 5% of the net assets tracking the benchmark index.[21] Second, OBX index weights are calculated based on market capitalization, a variable with significant positive

[20]Similarly, the use of so-called exchange traded funds (ETFs) has surged over the sample period (Ben-David et al. 2014). ETFs, like index funds, facilitate exposure towards, among other assets, baskets of stocks such as the OBX index. This surge, however, seems unlikely to explain my results. Although the literature is not conclusive, recent empirical evidence by Hamm (2014) suggests that ETF trading negatively correlates with measures of underlying stock liquidity. The driving mechanism, according to Hamm (2014), is that uninformed traders reduce their participation in the underlying asset if given the option to invest in ETFs, which reduces the liquidity of the underlying asset. If so, a surge in the ETF trading of OBX listed stocks would lead to opposite effects (from what I document).

[21]These figures are based on data from Morningstar in the time period June 2008 - April 2010. Net asset values are reported on different frequencies (monthly, quarterly, yearly) for different funds. Quarterly and yearly holdings are carried forward to create a monthly time-series. Average combined monthly net assets for funds tracking OBX in the sample period are approximately 5% of the net assets tracking OSEBX.

skewness (see Table 3). Consequently, for the marginal OBX stock, where the regression discontinuity effect is measured, this translates into a negligible index weight. For example, in the period $2008 - 2010$, the average index weight of the marginal OBX stock is 1.04%, which seems unlikely to explain the effects in Section 5.1. Finally, future OBX constituents are announced approximately two weeks prior to actual index reconstitution. Any excess trading activity caused by index benchmarking strategies is likely to be exhausted before the sample data begins, which is at revision date.

6.2 Control variables

If the RD design is valid, there is no need to control for observable characteristics (Lee and Lemieux 2010). This is because the randomness of treatment assignment close to the treatment threshold ensures that marginally included and excluded stocks, on average, are similar in their observable characteristics. Including control variables, however, may increase precision or even reduce estimation bias if observables are not entirely balanced between just-included and just-excluded stocks. In column four of Table 5, I add a comprehensive vector of firm characteristics to specification 4. Estimates of the effect of trader anonymity on stock liquidity and trading volume become slightly smaller in the inclusion of control variables but remain highly significant, both statistically and economically.

6.3 Confounding market structure trends ($2008 - 2010$)

European market structure developments unrelated to trader anonymity at the Oslo Stock Exchange but correlated with OBX index membership, could drive the results in Section 5.1. For example, the introduction of trader anonymity in 2008 coincides with the most disruptive market structure development in recent European equity trading history. Effective in late 2007, a new pan-European legislative (MiFID) abolished local stock exchange monopolies, and opened competition between exchanges. Anecdotal evidence suggests that entrant exchanges systematically challenged market shares in the most liquid shares before gradually expanding their selection of stocks.[22] Competition for order flow can correlate

[22]Multi-lateral trading facilities (MTFs) began competing for order flow in the largest OSE stocks first, then gradually expanded their selection. For example, the MTF Chi-X opened trading in the five largest OSE stocks in 2008 (Norsk Hydro ASA, Renewable Energy Corp A/S, StatoilHydro ASA, Telenor ASA, and Yara International ASA). Chi-X now offers trading in more than 50 OSE products. Similarly, the MTF

with OBX constituency, by virtue of being the most liquid shares at OSE, and confound the estimated effect of trader anonymity.

To address this concern, I generate a variable $Frag_{ie}$, which measures stock-level order flow fragmentation as the share of traded volume on competing trading venues relative to total traded volume across all venues (see appendix A.5 for details), and include it as a regressor in the baseline regression discontinuity design.[23] Column six of Table 5 shows that the estimates are robust to the inclusion of $Frag_{ie}$ as a regressor, which suggests that order flow fragmentation is not driving the results in Section 5.1.

Meanwhile, I am unable rule to out confounding effects from other concurrent market structure developments. The trader anonymity sample period $(2008-2010)$ is characterized by, among other things, an explosion in high-frequency trading (e.g. Jørgensen et al. 2014, Angel et al. 2011,2013), aggressive use (by stock exchanges) of new fee structures, such as maker-taker fees (e.g. Malinova and Park 2015), and a financial crisis. If these developments systematically correlate with OBX list membership, they may bias my estimates. To minimize the potential for time-varying confounders, in appendix A.6 I estimate a short-run difference-in-differences (DiD) model surrounding only the first introduction of trader anonymity at OSE. The DiD specification in appendix A.6 produces broadly the same results as the regression discontinuity design.

7 Mechanisms

In Section 5, I show that trader anonymity improves stock liquidity and trading volume. These results are consistent with theoretical models where informed traders supply and improve liquidity in anonymous markets, such as Boulatov and George (2013) and Rindi (2008).[24] To further explore the empirical support for these models, this section tests

Turquoise initially offered trading in 28 OSE stocks but has since expanded to offer trading in 169 OSE products.

[23]I include $Frag_{ie}$ in the baseline specification (equation 3), and not the extended RD model (equation 4), because $Frag_{ie} = 0$ for the entire period $2002-2007$. Including $Frag_{ie}$ in the extended RD model produces the same results.

[24]In Rindi (2008), the net effect of trader anonymity on stock liquidity depends on the exact modeling of information acquisition. When information acquisition is endogenous, anonymity improves liquidity, but when information acquisition is exogenous, anonymity degrades liquidity. In Boulatov and George (2013), the impact of anonymity on stock liquidity also depends on the aggressiveness by which informed traders supply liquidity. In their model, anonymity induces informed traders to supply rather than to demand liquidity

whether anonymity induces informed traders to transact more. I test this hypothesis using transaction-level data that allow me to create empirical proxies for informed and uninformed traders. Consistent with informed traders supplying and improving liquidity under anonymity, I find that the increase in aggregate trading volume documented in Section 5 is mostly accounted for by informed investors.

7.1 Data and descriptives

I use proprietary data on the trading of all investors at the OSE, obtained from the OSE. Each entry in this dataset is a trade and identifies the buying and selling brokerage firm as well as volumes, prices and stock identifiers (see Section 3 for more details). As empirical proxies for 'informed' and 'uninformed' investors, I follow Linnainmaa and Saar (2012) and use institutional and retail investors, respectively. Based on the brokerage firm identifiers reported in the data, I classify order flow from online discount brokerages as retail. The residual order flow is collectively referred to as 'institutional.' Similar to Linnainmaa and Saar (2012), I distinguish between foreign and domestic institutions. Appendix A.7 provides further detail on this order flow decomposition.

Table 6 describes the trading behavior of retail and institutional investors. Domestic institutions are the most active investors at the OSE with an average (stock-day) trading volume of 23 million NOK spread across 316 trades, followed by foreign institutions (13 million NOK, 215 trades) and finally retail investors (5 million NOK, 112 trades). To provide some evidence supporting that institutions are more sophisticated or 'informed' than retail investors, I follow Malinova et al. (2013) and compute intraday trading profits for each of the trader groups. The average per-stock-day trading loss of retail investors is 5.91 basis points.[25] Both the foreign and domestic institutions in my sample, in contrast, are able to generate positive trading profits, which is suggestive of higher sophistication among these traders. Moreover, consistent with previous literature comparing the trading behavior of

and, in addition, increases the aggressiveness by which they supply liquidity. The interaction of these effects generate improvements in stock liquidity under anonymity. In a recent theoretical framework, Roşu (2016) shows that when informed traders can choose whether to supply or demand liquidity, an exogenous increase in the share of informed traders in the market improves both stock liquidity and price efficiency. Roşu (2016), however, does not model the consequences of anonymous and non-anonymous markets.

[25]For comparison, Malinova et al. (2013) report average daily trading losses of 5.1 basis points for their sample of retail investors.

retail and institutional investors (e.g. Lee and Radhakrishna 2000, Barber et al. 2009), retail investors in my sample execute the smallest trades with an average value of 42 663NOK, while institutional trades average 47 824NOK (foreign) and 60 948NOK (domestic) in size.

7.2 Results

To explore whether anonymity induces informed traders to transact more, I estimate the causal impact of trader anonymity on the trading activity (trades and trading volume) of institutional and retail investors using a regression discontinuity design. The regression discontinuity design compares how the same group of investors trade in two otherwise similar stocks — those just-included and just-excluded from the OBX index — where trading in one stock is anonymous and in the other it is not.

I implement the regression discontinuity design with the same two-stage least-squares approach used in Section 5 to measure the impact of trader anonymity on stock liquidity and trading volume. In the first stage regression, I relate predicted treatment status $T_{ie} = \mathbf{1}[r_{ie} \geq 0]$ for stock i during event e to actual treatment status D_{ie}:

$$D_{ie} = \alpha_0 + \phi r_{ie} + \psi T_{ie} + \omega T_{ie} \times r_{ie} + \varpi_{ie}. \tag{5}$$

In the second stage regression, I relate the trading activity y^g_{ie} of trader group g to treatment status D_{ie}:

$$y^g_{ie} = \alpha_1 + \nu r_{ie} + \tau D_{ie} + \beta D_{ie} \times r_{ie} + \varepsilon_{ie}. \tag{6}$$

The inclusion of the ranking variable r_{ie}, which is centered on zero, in both regression stages ensures that all identification is centered on $r_{ie} = 0$ — the cutoff point between anonymous and transparent trading. Thus, the coefficient τ measures a discontinuous change in the trading behavior of investors exactly at $r_{ie} = 0$, properly scaled by the first stage relationship.

Figure 5 presents graphical evidence on the change in investor behavior at $r_{ie} = 0$. Table 7 presents estimates of τ using a bandwidth specification of $h = 15$. Table 7 shows that both foreign and domestic institutions transact much more frequently when they can do so anonymously. For example, foreign institutions increase their trading volumes by more than 300%, and domestic institutions more than double their trading volumes. In fact, since retail investors do not change their trading behavior in response to anonymity, the entire increase in aggregate trading volume documented in Section 5, can be attributed to institutions.

7.3 Discussion

I have shown that trader anonymity increases stock liquidity and trading volume, and that the increase in aggregate trading volume is mostly accounted for by increased activity by institutional investors. I interpret the simultaneous increase in institutional trading and stock liquidity under anonymity as consistent with theoretical frameworks where informed investors supply and improve liquidity in anonymous markets (e.g. Boulatov and George 2013, Rindi 2008).

Meanwhile, I cannot exclude the possibility that the positive relationship between institutional trading and stock liquidity is spurious.[26] This is because the data available do not allow me to detect whether investors supply liquidity (place limit orders) or demand liquidity (place market orders). Empirical evidence from other markets, however, suggest that informed traders causally improve stock liquidity through liquidity provision. For example, in experimental securities markets, Perotti and Rindi (2006) show that anonymous markets attract informed traders who supply and improve liquidity while Bloomfield et al. (2005) provide evidence that informed traders endogenously take on the role as liquidity suppliers. Similarly, Kaniel and Liu (2006) provide empirical evidence that liquidity providers at the NYSE are informed.

The positive relationship between institutional trading and stock liquidity observed in the current paper may also be explained by the order anticipation framework promoted by Friederich and Payne (2014). They argue that liquidity-motivated institutions, who are not necessarily informed, enter anonymous markets to avoid the transaction costs associated with exposing their trading demands in transparent markets. Because trader anonymity pro-

[26]Another possibility is that causality runs from stock liquidity to institutional trading, and not the other way around. For example, Collin-Dufresne and Fos (2015) present empirical evidence supporting a positive relationship between informed trading and stock liquidity and argue that it can be explained by two mechanisms — 1) informed traders strategically choose to trade when liquidity is high and 2) informed traders supply liquidity. By the latter mechanism, informed trading causally improves stock liquidity while by the former mechanism causality is reversed.

tects institutions from order anticipation (front-running), Friederich and Payne (2014) argue, anonymity allows institutions to patiently accumulate positions which adds competition to the market's liquidity supply — thus, improving stock liquidity. I interpret this empirical framework to be analogous to the theoretical informed liquidity supply framework promoted in the current paper since they both describe how trader anonymity induces a certain group of investors (who move prices under transparency) to transact and supply liquidity.

8 Conclusion

In this paper, I assess the effect of trader anonymity on stock liquidity and trading volume. An anonymity reform in the period 2008 − 2010 at the Oslo Stock Exchange provides a source of exogenous variation. The 25 most traded stocks at the OSE were semi-annually assigned to anonymity; all others were not. Comparing just-included and just-excluded stocks in a so-called regression discontinuity design provides causal estimates. I find that trader anonymity increases stock liquidity and trading volume. Retail investors, arguably the least informed investors in the market, do not adjust their trading behavior in response to anonymity. In contrast, institutional investors, perhaps the most informed market participants, transact much more when they can do so anonymously.

These results are consistent with theoretical models where informed traders supply and improve liquidity in anonymous markets (e.g. Boulatov and George 2013, Rindi 2008). The results are inconsistent with theoretical models that emphasize the adverse selection costs of anonymous markets — the main competing mechanism.

The results in this paper contribute to the existing empirical literature on anonymous trading in equities (e.g. Theissen 2003, Waisburd 2003, Comerton-Forde et al. 2005, Foucault et al. 2007, Thurlin 2009, Comerton-Forde and Tang 2009, Hachmeister and Schiereck 2010, Friederich and Payne 2014, Dennis and Sandås 2015). This literature is based on non-exogenous variation, where identification is difficult. In contrast, I exploit exogenous variation in trader anonymity for causal inference. My research contributes to this literature with clean identification and unambiguous results on the effect of anonymity on stock liquidity and trading volume.

My research may provide guidance to regulators in the United States and Europe who are currently considering comprehensive market structure changes to increase equity market transparency. My results suggest that increasing equity market transparency may in fact worsen overall market quality by discouraging informed traders from participating and providing liquidity to the market.

References

Ahern, K. R. and A. K. Dittmar (2012): "The Changing of the Boards: The Impact on Firm Valuation of Mandated Female Board Representation," *The Quarterly Journal of Economics*, 127, 137–197.

Angel, J., L. Harris, and C. Spatt (2011): "Equity Trading in the 21st Century," *Quarterly Journal of Finance*, 1, 1–53.

——— (2013): "Equity Trading in the 21st Century: An Update," .

Barber, B. and T. Odean (2008): "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors," *Review of Financial Studies*, 20, 785–818.

Barber, B. M., T. Odean, and N. Zhu (2009): "Do retail trades move markets?" *Review of Financial Studies*, 22, 151–186.

Ben-David, I., F. Franzoni, and R. Moussawi (2014): "Do ETFs Increase Volatility?" Working Paper 20071, National Bureau of Economic Research.

Benhami, K. (2006): "Liquidity provider's valuation of anonymity: The Nasdaq Market Makers evidence," Working Paper, Universite de Toulouse.

Bessembinder, H. (2003): "Trade Execution Costs and Market Quality after Decimalization," *Journal of Financial and Quantitative Analysis*, 28, 747–777.

Bloomfield, R. J., M. O'Hara, and G. Saar (2005): "The "make or take" decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics*, 75, 165 – 199.

Boone, A. L. and J. T. White (2015): "The effect of institutional ownership on firm transparency and information production," *Journal of Financial Economics*, 117, 508 – 533.

Boulatov, A. and T. J. George (2013): "Hidden and Displayed Liquidity in Securities Markets with Informed Liquidity Providers," *Review of Financial Studies*.

Buti, S., F. Consonni., B. Rindi., Y. Wen, and I. M. Werner (2015): "Tick Size: Theory and Evidence," Working Paper. Rotman School of Management Working Paper No. 2485069; Fisher College of Business Working Paper 2015-03-04; Charles A. Dice Center Working Paper No. 2015-04. Available at SSRN: http://ssrn.com/abstract=2485069.

Card, D. and D. Lee (2008): "Regression Discontinuity Inference With Specification Error," *Journal of Econometrics*, 142, 655–674.

Cellini, S. R., F. Ferreira, and J. Rothstein (2010): "The Value of School Facility Investments: Evidence from a Dynamic Regression Discontinuity Design," *Quarterly Journal of Economics*, 125, 215–261.

Chang, Y.-C., H. Hong, and I. Liskovich (2014): "Regression Discontinuity and the Price Effects of Stock Market Indexing," *Review of Financial Studies*.

Collin-Dufresne, P. and V. Fos (2015): "Do Prices Reveal the Presence of Informed Trading?" *The Journal of Finance*, 70, 1555–1582.

Comerton-Forde, C., A. Frino, and V. Mollica (2005): "The impact of limit order anonymity on liquidity: Evidence from Paris, Tokyo, and Korea," *Journal of Economics and Business*, 57, 528–540.

Comerton-Forde, C., T. J. Putnins, and K. M. Tang (2011): "Why Do Traders Choose to Trade Anonymously?" *Journal of Financial and Quantitative Analysis*, 46, 1025–1049.

Comerton-Forde, C. and K. Tang (2009): "Anonymity, Liquidity, and Fragmentation," *Journal of Financial Markets*, 12, 337–367.

Commission, E. (2014): "More transparent and safer financial markets: European Commission welcomes European Parliament vote on updated rules for Markets in Financial Instruments (MiFID II)," Available at: http://europa.eu/rapid/press-release_STATEMENT-14-129_en.htm.

Cuñat, V., M. Gine, and M. Guadalupe (2012): "The Vote Is Cast: The Effect of Corporate Governance on Shareholder Value," *The Journal of Finance*, 67, 1943–1977.

Dennis, P. J. and P. Sandås (2015): "Does Trading Anonymously Enchange Liquidity?" Working Paper, McIntire School of Commerce, University of Virginia.

Døskeland, T. M. and H. K. Hvide (2011): "Do Individual Investors Have Asymmetric Information Based on Work Experience?" *The Journal of Finance*, 66, 1011–1041.

Federal Register (2003): "Securities and Exchange Commission: Self-Regulatory Organizations; Notice of Filing of Proposed Rule Change by the National Association of Securities Dealers, Inc. and Amendments No. 1 and 2 Thereto Relating to a Post-Trade Anonymity Feature in SuperMontage," *Federal Register*, 68, 39605.

FINRA (2016): "16-14 FINRA Announces Implementation Date for Publication of ATS Block-Size Trade Data; Implementation Date: October 3, 2016," Available at: http://finra.complinet.com/en/display/display.html?rbid=2403&element_id=12304.

Foucault, T., S. Moinas, and E. Theissen (2007): "Does Anonymity Matter in Electronic Limit Order Markets," *Review of Financial Studies*, 20, 1707–1747.

Foucault, T., M. Pagano, and A. Röell (2013): *Market Liquidity: Theory, Evidence, and Policy*, Oxford University Press.

Friederich, S. and R. Payne (2014): "Trading anonymity and order anticipation," *Journal of Financial Markets*, 21, 1 – 24.

Hachmeister, A. and D. Schiereck (2010): "Dancing in the dark: post-trade anonymity, liquidity and informed trading," *Review of Quantitative Finance and Accounting*, 34, 145–177.

Hahn, J., P. Todd, and W. van der Klaauw (2001): "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, 69, 201–209.

Hamm, S. J. W. (2014): "The Effect of ETFs on Stock Liquidity," Working Paper. Available at SSRN: http://ssrn.com/abstract=1687914.

Harris, L. E. and E. Gurel (1986): "Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures," *Journal of Finance*, 41, 815–829.

Hirshleifer, D., S. S. Lim, and S. H. Teoh (2009): "Driven to distraction: Extraneous events and underreaction to earnings news," *Journal of Finance*, 64, 2289–2325.

Hope, B. (2013): "'Upstairs' Trading Draws More Big Investors," *The Wall Street Journal*, available at: http://www.wsj.com/articles/SB10001424052702303722104579240393018918698.

Huddart, S., J. Hughes, and C. Levine (2001): "Public Disclosure and Dissimulation of Insider Trades," *Econometrica*, 69, 665–681.

Jørgensen, K., J. Skjeltorp, and B. A. Ødegaard (2014): "Throttling hyperactive robots - Message to trade ratios at the Oslo Stock Exchange," Working Paper, University of Stavanger.

Kaniel, R. and H. Liu (2006): "So what orders to informed traders use?" *Journal of Business*, 79, 1867–1913.

Lee, C. M. C. and B. Radhakrishna (2000): "Inferring Investor Behavior: Evidence from TORQ Data," *Journal of Financial Markets*, 3, 83–111.

Lee, D. (2008): "Randomized Experiments from Non-random Selection in U.S. House Elections," *Journal of Econometrics*, 142, 675–697.

Lee, D. and T. Lemieux (2010): "Regression discontinuity designs in economics," *Journal of Economic Literature*, 48, 281–355.

Linnainmaa, J. T. and G. Saar (2012): "Lack of anonymity and the inference from order flow," *Review of Financial Studies*, 25, 1414–1456.

Malinova, K. and A. Park (2015): "Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality," *Journal of Finance*, 70, 509–536.

Malinova, K., A. Park, and R. Riordan (2013): "Shifting Sands: High Frequency, Retail, and Institutional Trading Profits over Time," Working Paper, Available at SSRN: http://ssrn.com/abstract=2183806.

Næs, R., J. Skjeltorp, and B. A. Ødegaard (2011): "Stock Market Liquidity and The Business Cycle," *Journal of Finance*, LXVI, 139–176.

PEROTTI, P. AND B. RINDI (2006): "Market for Information and Identity Disclosure in an Experimental Open Limit Order Book," *Economic Notes*, 35, 97–119.

PHAM, T. P., P. L. SWAN, AND P. J. WESTERHOLM (2014): "Shining a Spotlight on Counterparty Identity in the World's Best-Lit Market," Available at SSRN: http://ssrn.com/abstract=2644149.

RINDI, B. (2008): "Informed Traders as Liquidity Providers: Anonymity Liquidity and Price Formation," *Review of Finance*, 12, 497–532.

ROŞU, I. (2016): "Liquidity and Information in Order Driven Markets," Working Paper, HEC Paris.

SHLEIFER, A. (1986): "Do Demand Curves for Stocks Slope Down?" *Journal of Finance*, 59, 579–590.

THEISSEN, E. (2003): "Trader Anonymity, Price Formation and Liquidity," *European Finance Review*, 7, 1–26.

THURLIN, A. (2009): "Pre-trade Transparency, Market Quality, and Informed Trading," *Working Paper*, Hanken School of Economics.

U.S. SECURITIES AND EXCHANGE COMMISSION (2009): "Speech by SEC Chairman: Statement on Dark Pool Regulation Before the Commission Open Meeting," Available at: https://www.sec.gov/news/speech/2009/spch102109mls.htm.

VAN KERVEL, V. AND A. J. MENKVELD (2015): "High-Frequency Trading around Large Institutional Orders," Working Paper, VU University Amsterdam.

WAISBURD, A. C. (2003): "Anonymity and liquidity: evidence from the Paris Bourse," Working Paper, Neeley School of Business.

WORLD FEDERATION OF STOCK EXCHANGES (2016): "Statistics: Monthly Reports," Data available at: https://www.world-exchanges.org/home/index.php/statistics/monthly-reports.

9 Tables

Table 1: OBX revisions 2002 − 2010

Treatment revisions (June 2008 - April 2010)			
Event	*Announced*	*Revision*	*Turnover period*
4	9 Dec 2009	18 Dec 2009	1 June 2009 - 30 Nov 2009
3	11 June 2009	19 June 2009	1 Dec 2008 - 29 May 2009
2	11 Dec 2008	19 Dec 2008	1 June 2008 - 30 Nov 2008
1	9 June 2008	20 June 2008	1 Dec 2007 - 31 May 2008
Placebo revisions (June 2002 - December 2007)			
Event	*Announced*	*Revision*	*Turnover period*
0	07 Dec 2007	21 Dec 2007	1 June 2007 - 30 Nov 2007
-1	13 June 2007	22 June 2007	1 Dec 2006 - 31 May 2007
-2	11 Dec 2006	22 Dec 2006	1 June 2006 - 30 Nov 2006
-3	12 June 2006	16 June 2006	1 Dec 2005 - 31 May 2006
-4	12 Dec 2005	16 Dec 2005	1 June 2005 - 30 Nov 2005
-5	07 June 2005	17 June 2005	1 Dec 2004 - 31 May 2005
-6	10 Dec 2004	17 Dec 2004	1 June 2004 - 30 Nov 2004
-7	10 June 2004	18 June 2004	1 Dec 2003 - 31 May 2004
-8	12 Dec 2003	19 Dec 2003	1 June 2003 - 30 Nov 2003
-9	12 June 2003	20 June 2003	1 Dec 2002 - 31 May 2003
-10	10 Dec 2002	20 Dec 2002	1 June 2002 - 30 Nov 2002
-11	13 June 2002	21 June 2002	1 Dec 2001 - 31 May 2002

Note: The table presents announcement dates and revision dates for all OBX list revisions in the time period 2002 − 2010. On OBX revision dates (*Revision*), the 25 stocks with the highest currency trading volume accumulated over the previous six months (*Turnover period*) are chosen from the broader index OSEBX to comprise the OBX the subsequent six months. New OBX stock compositions are announced 1-2 weeks before revisions. Revisions of the OBX list between June 2, 2008 and April 12, 2010, assigned trader anonymity to OBX listed stocks. Revisions of the OBX list before this period did not assign trader anonymity to OBX listed stocks.

Table 2: Examples of anonymous and non-anonymous trade feeds

Non-anonymous trade feed				
Broker ID (buy)	*Broker ID (sell)*	*Stock ID*	*Volume*	*Price*
ESO	NEO	STL	500	195.6
NON	NON	NEC	4000	8.13
ND	ND	QEC	2000	20.9
JPM	NEO	DNBNOR	600	71.9
UBS	NEO	DNBNOR	1600	71.9
ESO	NEO	DNBNOR	700	71.9
NON	NEO	DNBNOR	1400	71.9
LBI	SHB	EKO	1200	81.5
Anonymous trade feed				
Broker ID (buy)	*Broker ID (sell)*	*Stock ID*	*Volume*	*Price*
.	.	STL	500	195.6
.	.	NEC	4000	8.13
.	.	QEC	2000	20.9
.	.	DNBNOR	600	71.9
.	.	DNBNOR	1600	71.9
.	.	DNBNOR	700	71.9
.	.	DNBNOR	1400	71.9
.	.	EKO	1200	81.5

Note: The table illustrates the difference between post-trade anonymity and post-trade non-anonymity. The top panel shows the information available to market participants when trading is non-anonymous. The bottom panel shows the information available to market participants when trading is anonymous.

Table 3: Summary statistics

	Sample descriptives (2002–2010)					
	μ	σ	Min.	Median	Max.	N
Trading characteristics						
Share volume	1.6	6.6	0.0	0.2	469.3	72201
Currency volume	81.3	251.0	0.0	7.5	10345.0	72201
Trades	451.2	930.1	0.0	88.0	19510.0	72298
Trade value	154.3	562.1	0.3	80.1	59731.2	69278
Trade size	4326.7	20921.2	5.7	1601.1	2370560.0	69278
Relative spread (bps.)	148.8	318.2	0.6	66.9	14482.8	72097
Firm fundamentals						
Market cap.	18628.2	50919.5	46.5	4484.2	538881.4	71114
Total equity	8934.3	21397.3	-859.1	1943.8	214079.0	71747
Total assets	40189.1	156162.5	0.0	5663.9	1832699.0	71747
Price/Book	3.8	5.8	-10.7	2.3	60.7	70627
Operating profit	3266.0	17627.1	-14574.0	343.5	228858.6	69789
Operating income	17515.5	57054.5	0.1	2670.7	651977.0	69383
Cash and deposits	1405.7	2997.3	0.9	452.0	27148.0	69090
	OBX vs Non-OBX (2002–2010)					
	μ^{OBX}	$\mu^{Non\,OBX}$	Diff.	$\sigma^{diff.}$	N_1	N_2
Trading characteristics						
Share volume	3.1	0.6	2.5***	0.0	27992	44209
Currency volume	194.2	9.9	184.3***	1.8	27992	44209
Trades	1020.5	91.5	928.9***	6.2	27995	44303
Trade value	184.2	134.0	50.2***	4.3	27987	41291
Trade size	3131.3	5136.9	-2005.6***	161.8	27987	41291
Relative spread (bps.)	44.4	215.2	-170.8***	2.3	27992	44105
Firm fundamentals						
Market cap.	41250.0	3941.1	37308.9***	364.9	27995	43119
Total equity	18840.9	2708.0	16132.9***	152.6	27690	44057
Total assets	91600.9	7876.5	83724.3***	1156.1	27690	44057
Price/Book	3.5	4.0	-0.5***	0.0	26570	44057
Operating profit	6935.5	1113.6	5822.0***	136.5	25802	43987
Operating income	38431.3	5395.4	33035.9***	431.6	25455	43928
Cash and deposits	2896.2	532.4	2363.8***	21.8	25527	43563

Note: The top panel lists means (μ), standard deviations (σ), minimum (Min.) and maximum values (Max.), medians, and number of observations (N) for the full sample of stock-day observations, the first 70 trading days after each OBX revision in the period 2002-2010. Share volume is the number of shares traded, in million shares. Currency volume is the value of shares traded, in million NOK. Trades is the number of transactions. Trade value is currency volume divided by trades, expressed in thousand NOK. Trade size is share volume divided by trades. All stock fundamendals, except Price/Book, are expressed in million NOK. The bottom table shows a t-test of different means between OBX index stocks, and non-OBX stocks, for all observations in the period 2002-2010. μ^{OBX} and μ^{NonOBX} represent the means for OBX and non-OBX stocks, respectively. Diff. is the difference between μ^{OBX} and μ^{NonOBX}. $\sigma^{diff.}$ is the standard error of the difference-in-means. N_1 is the number of observations in the OBX sample. N_2 is the number of observations in the non-OBX sample.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table 4: Main results

	Bandwidth		
	h=10	*h=15*	*h=20*
Dep. variable: Relative spread (log)			
τ	-0.86***	-0.56***	-0.56***
	(-4.33)	(-4.41)	(-6.03)
%Δ	-57.82	-42.85	-42.77
N	80	120	160
Dep. variable: Trades (log)			
τ	0.99***	0.62**	0.51***
	(3.27)	(2.63)	(3.35)
%Δ	169.22	85.12	67.24
N	80	120	160
Dep. variable: Trading volume (log)			
τ	1.25***	0.63**	0.45**
	(3.12)	(2.25)	(2.47)
%Δ	247.49	88.62	57.36
N	80	120	160
Dep. variable: Volatility			
τ	-0.00	-0.00	0.00
	(-0.02)	(-0.31)	(0.82)
%Δ	-	-	-
N	80	120	160

Note: The table gives estimates of τ from the baseline fuzzy regression discontinuity design (equation 3). Relative spread is the end-of-day quoted spread divided by the end-of-day quote midpoint, log-transformed. Trades is the daily number of transactions, log-transformed. Value traded is the daily currency trading volume, log-transformed. These outcomes are first defined on the stock-day level, and subsequently averaged into a single stock-event observation. Volatility is the variance of the close-to-close returns process, computed at the stock-event level. τ is estimated in a two-stage least-squares (2SLS) specification, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. $T_{ie} = \mathbf{1}\left[r_{ie} \geq 0\right]$, where r_{ie} is a ranking variable determined each June and December in $2008 - 2010$, by previous six months trading volume. r_{ie} has been normalized to zero by subtracting it from 25. $D_{ie} \times r_{ie}$ is included in the estimation to allow r_{ie} to vary with D_{ie}, and is instrumented with $T_{ie} \times r_{ie}$. Exogenous controls include the ranking variable r_{ie} and $ticksize_{ie}$. The 2SLS is estimated non-parametrically within bandwidths h. h indicates the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). Percentage change, %Δ, is calculated as $e^{\tau}-1$. Standard errors are clustered at r_{ie}. t-statistics in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Robustness specifications

	Robustness specifications					
	1	*2*	*3*	*4*	*5*	*6*
Dep. variable: Relative spread (log)						
τ	-0.56***	0.11	0.07	0.15	0.10	-0.45***
	(-4.41)	(0.61)	(0.64)	(0.78)	(0.68)	(-2.88)
τ^{Diff}			-0.60***	-0.62***	-0.37***	
			(-6.81)	(-6.12)	(-3.68)	
%Δ	-42.85	11.36	-44.92	-46.01	-30.66	-36.55
N	120	360	480	480	461	111
Dep. variable: Trades (log)						
τ	0.62**	0.03	-0.05	-0.01	0.03	0.60**
	(2.63)	(0.07)	(-0.19)	(-0.03)	(0.09)	(2.30)
τ^{Diff}			0.79***	0.79***	0.56***	
			(4.36)	(3.86)	(2.79)	
%Δ	85.12	2.61	121.23	119.70	74.77	81.50
N	120	360	480	480	461	111
Dep. variable: Trading volume (log)						
τ	0.63**	0.02	-0.07	-0.10	0.02	0.47
	(2.25)	(0.05)	(-0.19)	(-0.19)	(0.04)	(1.29)
τ^{Diff}			0.79***	0.81***	0.43*	
			(3.57)	(3.16)	(1.83)	
%Δ	88.62	2.45	121.14	124.98	53.06	60.48
N	120	360	480	480	461	111
Dep. variable: Volatility						
τ	-0.00	0.00	0.00	0.00	0.00	0.00
	(-0.31)	(0.79)	(0.44)	(0.66)	(0.36)	(1.05)
τ^{Diff}			0.00	-0.00	0.00	
			(0.00)	(-0.13)	(1.30)	
%Δ	-	-	-	-	-	-
N	120	360	480	480	461	111
Placebo	No	Yes	No	No	No	No
Pre-Post	No	No	Yes	Yes	Yes	No
Mover dummies	No	No	No	Yes	Yes	No
Controls	No	No	No	No	Yes	No
Fragmentation	No	No	No	No	No	Yes

Note: The table gives estimates of τ from extensions of the fuzzy regression discontinuity design (equation 3). The baseline specification is a two-stage least-squares (2SLS) approach, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. $T_{ie} = \mathbf{1}[r_{ie} \geq 0]$, where r_{ie} is a ranking variable determined each June and December in $2002 - 2010$, by previous six months trading volume. r_{ie} has been normalized to zero by subtracting it from 25. $D_{ie} \times r_{ie}$ is included in the estimation to allow r_{ie} to vary with D_{ie}, and is instrumented with $T_{ie} \times r_{ie}$. Exogenous controls include the ranking variable r_{ie} and $ticksize_{ie}$. The 2SLS is estimated within a bandwidth $h = 15$. h indicates the number of stocks included on either side of the treatment cutoff $r_{ie} = 0$. Column one gives the baseline specification, using data from $2008 - 2010$. Column two applies the baseline specification to the period $2002 - 2007$. Column three estimates the difference-in-differences between treatment periods ($2008 - 2010$) and placebo periods ($2002 - 2007$). Column four adds to the difference-in-differences model separate dummy variables for stocks moving in and out of the OBX list. Column five adds to column four a set of control variables (market capitalization (log), stock price (log), price-to-book (log), shares issued (log), total equity (log), total assets (log), operating profit, operating income, and cash and deposits (log)). Column six adds to the baseline specification a proxy for order flow fragmentation. This proxy is defined as the share of currency trading volume that occurs on all trading venues (dark pools, lit order books, SIs, and other OTC) excluding OSE, relative to the total currency trading volume across all trading venues including OSE. Percentage change, %Δ, is calculated as $e^{\tau}-1$. Standard errors are clustered at r_{ie}. t-statistics in parentheses.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table 6: Summary statistics by trader group

	μ	σ	Min.	Median	Max.	N
Retail						
Intraday returns	-5.91	61.54	-207.75	-1.49	219.63	1901
Currency volume	5.62	9.75	0.00	2.03	100.10	1874
Share volume	444.37	1154.89	0.03	45.71	16116.30	1874
Price paid	58.25	62.34	1.01	36.90	371.17	1874
Trades	112.81	169.58	1.00	51.00	1566.00	1874
Trade value	42663.09	25056.54	576.50	37719.53	169587.50	1874
Institutions^f						
Intraday returns	0.84	63.65	-207.75	-1.45	219.63	1895
Currency volume	12.95	23.77	0.00	5.01	365.06	1872
Share volume	465.81	1016.11	0.00	124.28	10702.63	1872
Price paid	58.68	62.45	1.01	36.91	368.86	1872
Trades	215.66	244.67	1.00	122.00	2443.00	1872
Trade value	47824.32	46988.93	63.00	40605.51	856272.13	1872
Institutions^d						
Intraday returns	3.37	42.72	-207.75	1.50	219.63	1910
Currency volume	23.16	40.52	0.00	10.85	792.38	1897
Share volume	1094.15	2684.03	0.26	267.40	32146.63	1897
Price paid	58.18	62.15	1.01	36.83	368.63	1897
Trades	316.64	335.94	1.00	219.00	2527.00	1897
Trade value	60948.23	53079.29	3221.67	50682.63	1204010.50	1897

Note: The table provides summary statistics separately for the trading of retail investors, foreign institutions ($Institutions^f$), and domestic institutions ($Institutions^d$). Observations are at the stock-day-trader group level, aggregated from transaction-level data covering stocks ranked $r_{ie} \in [-15, 14]$, during the four trader anonymity events in the period $2008 - 2010$. Intraday returns are computed as $\frac{sell_{itg}^{Value} - buy_{itg}^{Value} + \left(buy_{itg}^{shares} - sell_{itg}^{shares}\right) \times ClosingPrice_{it}}{sell_{itg}^{Value} + buy_{itg}^{Value}}$, where $sell_{itg}^{Currency} - buy_{itg}^{Currency}$ is the profit from intraday trading. The term $buy_{itg}^{shares} - sell_{itg}^{shares}$ is the end-of-day position, assuming zero inventory at the beginning of each day, which is evaluated at the closing price. $sell_{itg}^{Currency} + buy_{itg}^{Currency}$ is the overall traded currency volume. Intraday returns are expressed in basis points, and are winsorized at the 1 per cent level. Currency volume is the daily total trading volume, in millions NOK. Share volume is the daily share volume, in thousand shares. Price paid is the daily average per-share price paid, in NOK. Trade value is the daily average transaction value, in NOK. In contrast to intraday returns, which are computed over both buy and sell transactions, Currency volume, Share volume, Price paid, Trades, and Trade value, are computed over buy transactions only, to avoid double-counting transactions. The table lists means (μ), standard deviations (σ), minimum (Min.) and maximum values (Max.), medians, and number of observations (N).

Table 7: Heterogeneity in trader response

	Trader group		
	Retail	*Institutions^f*	*Institutions^d*
Dep. variable: Trades (log)			
τ	0.08	1.11**	0.63**
	(0.14)	(2.43)	(2.72)
%Δ	8.42	202.43	87.13
N	120	120	120
Dep. variable: Trading volume (log)			
τ	0.16	1.45**	0.72**
	(0.24)	(2.40)	(2.10)
%Δ	17.02	326.69	105.82
N	120	120	120

Note: The table gives estimates of τ from the fuzzy regression discontinuity design (equation 3). The RD design is estimated separately for retail investors, foreign institutions ($Institutions^f$), and domestic institutions ($Institutions^d$). The outcomes considered are the daily number of trades (log) and daily monetary trading volume (log). These outcomes are first computed on the stock-day-trader group level, then averaged into a single stock-event-trader group observation. Trading volume and Trades are computed using buy transactions only, to avoid double-counting transactions. τ is estimated in a two-stage least-squares specification, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. $T_{ie} = \mathbf{1}[r_{ie} \geq 0]$, where r_{ie} is a ranking variable determined each June and December in $2008 - 2010$, by previous six months trading volume. r_{ie} has been normalized to zero by subtracting it from 25. $D_{ie} \times r_{ie}$ is included in the estimation to allow r_{ie} to vary with D_{ie}, and is instrumented with $T_{ie} \times r_{ie}$. Exogenous controls include the ranking variable r_{ie} and $ticksize_{ie}$. The 2SLS is estimated within a bandwidth $h = 15$. h indicates the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). Percentage change, %Δ, is calculated as $e^{\tau}-1$. Standard errors are clustered at r_{ie}. t-statistics in parentheses.
** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

10 Figures

Figure 1: Time-line

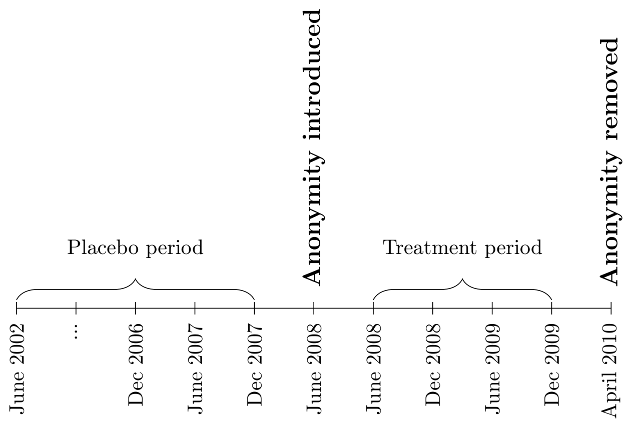

Note: The figure presents a time-line of the introduction and removal of anonymous trading at the Oslo Stock Exchange (OSE). Each tick on the time-line represents a revision of the OBX list composition. The OBX list is composed of the most traded stocks at OSE, and the composition is revised twice a year (June and December). OSE introduced post-trade anonymity of brokerages on June 2, 2008 for constituents of the OBX list. Anonymity was removed April 12, 2010. In the period June 2 to April 12 the OBX list was revised four times, each time giving anonymity to a new set of constituent stocks (Treatment period). The OBX list was also created before June 2, 2008 but constituents did not receive anonymity (Placebo period).

Figure 2: Illustration empirical design

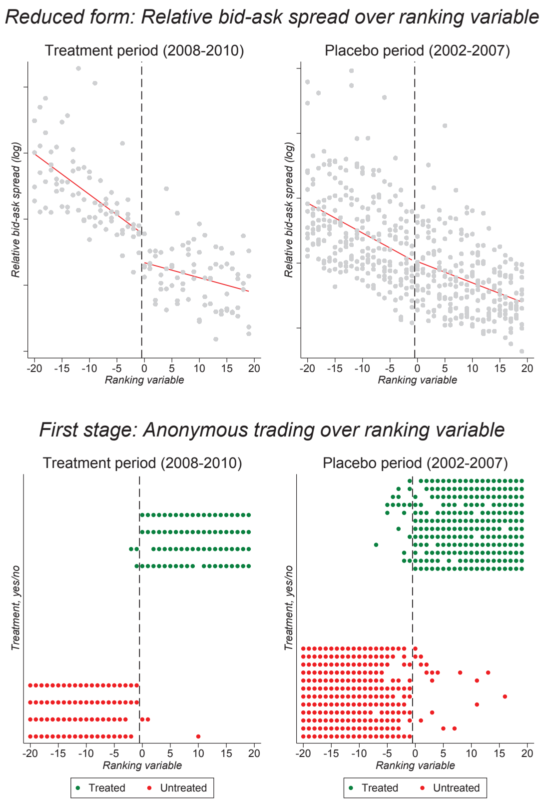

Note: The figure illustrates the fuzzy regression discontinuity (FRD) design applied to the logarithm of relative bid-ask spreads, a commonly used measure of illiquidity and transaction costs. The bottom panel relates treatment assignment to the ranking variable. Stocks are ranked semi-annually (June and December) based on previous six months trading activity. All observations to the right of the vertical break are intended for treatment based on this ranking variable. The ranking variable (r_{ie}) has been normalized to zero by subtracting it from 25. Green observations receive treatment, red observations do not. In the period 2008-2010, treatment implies anonymous trading and OBX index constituency. In the period 2002-2007, treatment implies OBX index constituency alone. The top panel relates relative bid-ask spreads to the same ranking variable. Linear regressions are fit separately on both sides of the vertical break. The FRD design estimates the effect of anonymous trading on relative bid-ask spreads as the vertical distance between regression intercepts at the vertical break, properly scaled by the treatment probability discontinuity at the same point.

Figure 3: Smoothness of stock characteristics

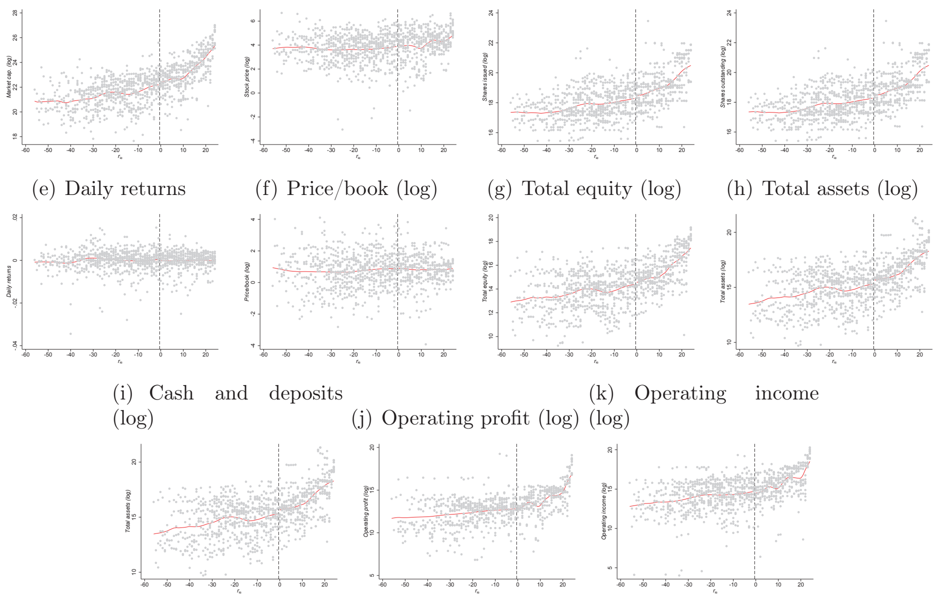

(a) Market cap. (log) (b) Stock price (log) (c) Shares issued (log) (d) Shares outstanding (log)

(e) Daily returns (f) Price/book (log) (g) Total equity (log) (h) Total assets (log)

(i) Cash and deposits (log) (j) Operating profit (log) (k) Operating income (log)

Note: The figure presents evidence on the smoothness of selected stock characteristics across the treatment threshold $r_{ie} = 0$ for all realizations of r_{ie} in the period $2002 - 2010$. The characteristics are market capitalization (log), stock price (log), shares issued (log), daily returns, price-to-book (log), total equity (log), operating profits (log), and operating income (log). The figure relates these characteristics to the ranking variable, r_{ie}, which is computed semi-annually (June and December) based on previous six months trading volume. Local polynomial regressions (red) are fit separately on both sides of the vertical break ($r_{ie} = 0$).

Figure 4: Coefficient estimates and bandwidth choice

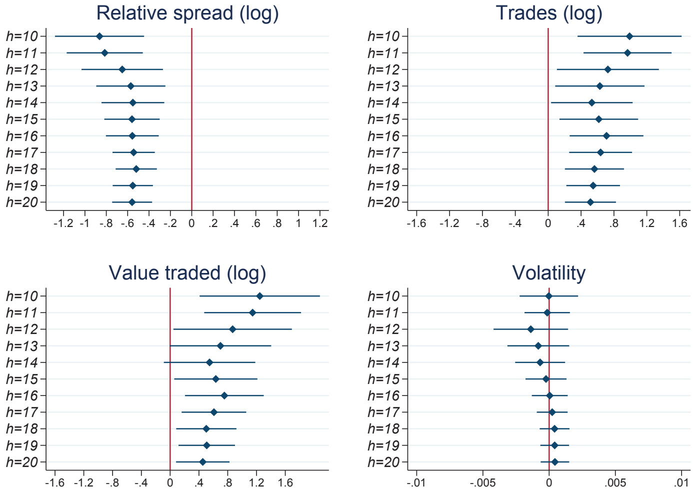

Note: The figure presents estimates of τ from the fuzzy regression discontinuity (RD) design (equation 3), with corresponding 95% confidence bands. Standard errors are corrected for clustering at r_{ie}, with a finite sample adjustment. The RD design is estimated non-parametrically within bandwidths h. h indicates the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). The figure presents estimates from $h \in [10, 20]$. τ is estimated in a two-stage least-squares (2SLS) specification, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. $T_{ie} = \mathbf{1}\left[r_{ie} \geq 0\right]$, where r_{ie} is a ranking variable determined in each June and December in $2008 - 2010$ by previous six months trading volume. r_{ie} has been normalized to zero by subtracting it from 25. $D_{ie} \times r_{ie}$ is included in the estimation to allow r_{ie} to vary with D_{ie}, and is instrumented with $T_{ie} \times r_{ie}$. Exogenous controls include the ranking variable r_{ie} and $ticksize_{ie}$.

Figure 5: Log(Trades) over r_{ie}, by trader group

(a) *Retail* (b) $Institutions^f$ (c) $Institutions^d$

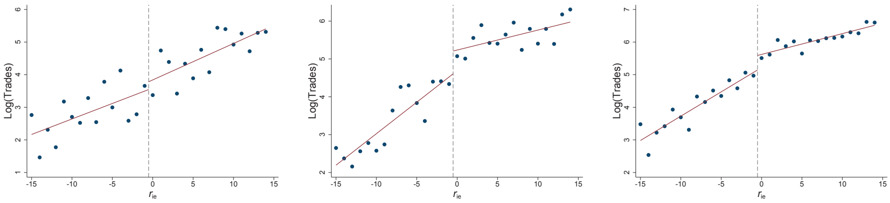

Note: The figure plots the natural logarithm of number of trades over the ranking variable r_{ie}, separately for retail investors, foreign institutions ($Institutions^f$), and domestic institutions ($Institutions^d$). Stocks to the right of the vertical break ($r_{ie} = 0$) are predicted for anonymous trading while stocks to the left of the vertical break are predicted for transparency. Log(Trades) is first computed on the stock-day-trader group level, then averaged into a single stock-event-trader group observation for each of the four realizations of r_{ie} in the period $2008 - 2010$. Each observation in the figure represents the average across these four realizations of r_{ie}.

A Appendix

A.1 Overview: Trader anonymity policy changes

The choice of transparency is one of the most hotly debated issues in equity market regulation, as it can affect price discovery, liquidity, and the distribution of rents between market participants (Foucault et al. 2013). Market transparency is defined by the amount of trading information, on prices, quantities, or identities, that is available to participants. Desirable transparency is determined by the individual trading venue and varies substantially between markets and over time.

Many leading stock exchanges, such as the Nasdaq, London Stock Exchange, and Deutsche Börse, have reduced transparency over the last decade by increasing trader anonymity. Practically, trader anonymity is implemented by concealing trader identifiers from orders in the order book (pre-trade anonymity) and/or from completed transactions in the trade feed (post-trade anonymity). Table A.1 summarizes recent stock exchange policy changes to both forms of trader anonymity. The summary focuses on trader anonymity policy changes analyzed in academic articles, or policy changes that have received attention by the media, and is not exhaustive.

Table A.1: Summary of trader anonymity policy changes

Exchange	**Event date**	**Policy change**	**Source**
Copenhagen	March 13, 2006	Introduction pre-trade anonymity	Nasdaq OMX note [a]
Frankfurt	March 27, 2003	Introduction post-trade anonymity	Hachmeister and Schiereck (2010)
Helsinki	March 13, 2006	Introduction pre-trade anonymity	Thurlin (2009)
Helsinki	June 2, 2008	Introduction post-trade anonymity	Dennis and Sandås (2015)
Helsinki	April 14, 2009	Removal post-trade anonymity	Nasdaq OMX note [b]
Istanbul	October 8, 2010	Introduction post-trade anonymity	ISE note [c]
London	February 26, 2001	Introduction post-trade anonymity	Friederich and Payne (2014)
Nasdaq	December, 2002	Increased pre-trade anonymity	Benhami (2006)
Nasdaq	October, 2003	Increased post-trade anonymity	Benhami (2006)
Oslo	October 22, 2007	Introduction pre-trade anonymity	OSE officials
Oslo	June 2, 2008	Introduction post-trade anonymity	This paper
Oslo	April 12, 2010	Removal post-trade anonymity	This paper
Paris	April 23, 2001	Introduction pre-trade anonymity	Foucault et al. (2007)
Reykjavik	June 2, 2008	Introduction post-trade anonymity	Dennis and Sandås (2015)
Riga	November 1, 2007	Introduction pre-trade anonymity	Nasdaq Baltic note [d]
Seoul	November 25, 1996	Removal post-trade anonymity	Pham et al. (2014)
Seoul	October 25, 1999	Removal pre-trade anonymity	Comerton-Forde et al. (2005)
Stockholm	March 13, 2006	Introduction pre-trade anonymity	Nasdaq OMX note [a]
Stockholm	June 2, 2008	Introduction post-trade anonymity	Dennis and Sandås (2015)
Stockholm	April 14, 2009	Removal post-trade anonymity	Nasdaq OMX note [b]
Sydney	November 28, 2005	Introduction pre-trade anonymity	Comerton-Forde and Tang (2009)
Tallinn	November 1, 2007	Introduction pre-trade anonymity	Nasdaq Baltic note [d]
Tokyo	June 30, 2003	Introduction pre-trade anonymity	Comerton-Forde et al. (2005)
Toronto	March 22, 2002	Introduced voluntary trader anonymity	Comerton-Forde et al. (2011)
Vilnius	November 1, 2007	Introduction pre-trade anonymity	Nasdaq Baltic note [d]

Note: The table gives an overview of stock exchange policy changes in trader anonymity. [a] *Changing the Nordic Market Microstructure*, April 2007.
[b] *NASDAQ OMX changes Post Trade Anonymity for the equity market trading in stockholm and Helsinki*, March 2009.
[c] *Markets and Operations*, October 2011.
[d] *Implementation of pre-trade anonymity*, November 2007.

A.2 First-stage regressions

In Section 5, I use a two-stage least-squares approach to estimate the causal effect of trader anonymity on stock outcomes. The specification in Section 5 is a fuzzy regression discontinuity (RD) design, where the predicted treatment T_{ie} (predicted by previous six month trading volume) is used as an instrumental variable for the actual treatment D_{ie}. In this section, I report the first-stage regressions of the fuzzy RD design. The first-stage regressions relate D_{ie} to T_{ie} and the ranking variable r_{ie}:

$$D_{ie} = b + \phi r_{ie} + \psi T_{ie} + \omega T_{ie} \times r_{ie} + \varphi ticksize_{ie} + \varpi_{ie}. \quad (7)$$

Estimates of ψ are presented in Table A.2. I present estimates separately for bandwidths $h = 10$, $h = 15$, and $h = 20$. The bandwidth is the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). Standard errors are clustered by r_{ie} with a finite sample adjustment; t-statistics are in parentheses.

ψ increases from a low of 0.55 at $h = 10$ to a high of 0.76 at $h = 20$. All point estimates are statistically significant, t-statistics increasing from 4.16 ($h = 10$) to 7.93 ($h = 20$). Crossing the treatment threshold increases the probability of treatment by 55% - 76%. The larger the bandwidth, the stronger the instrument. The centered R^2 varies around 0.80 for all bandwidths. R^2 is not necessarily monotonically increasing in h because of variation in $ticksize_{ie}$. The Angrist-Pischke multivariate F test of excluded instruments (statistic not reported in table) shows that T_{ie} is a sufficiently strong instrument for all bandwidths.

Table A.2: First-stage regressions of fuzzy RDD

	Bandwidth		
	h=10	*h=15*	*h=20*
ψ	0.55***	0.69***	0.76***
	(4.16)	(6.09)	(7.93)
R^2	0.80	0.82	0.87
F	216.92	287.85	812.66
N	80	120	160

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

A.3 Robustness tests: Polynomials in r_{ie}

In the main analysis, I investigate how trader anonymity at the Oslo Stock Exchange (OSE) affects market quality by using a regression discontinuity (RD) design. With the RD design, I find that anonymity increases stock liquidity (smaller relative bid-ask spreads) and trading activity (number of trades and traded value) but has no effect on returns volatility. Unbiased estimation of the RD effects requires an assumption about the functional form of the relationship between the running variable r_{ie} and outcomes y_{ie}. The RD literature has proposed two main approaches to estimating the RD design when this functional form is unknown. The first approach, which I use in the main text, is to estimate the RD design non-parametrically with so-called local linear regressions. This approach implies estimating linear regressions within a confined estimation range surrounding the treatment threshold $r_{ie} = 0$.

The second approach, which I take in this appendix, is to expand the estimation range and allow for a flexible relationship between y_{ie} and r_{ie} through a polynomial expansion in r_{ie}. The benefits of this approach are twofold. First, using a larger portion of the overall sample may increase the statistical precision of the estimation procedure. In this section, I use a bandwidth $h = 25$ (the number of stocks included on either side of $r_{ie} = 0$). This is the widest bandwidth attainable in my setting, while still preserving a symmetric sample surrounding $r_{ie} = 0$. Second, reporting estimates from a wider range of regression discontinuity specifications increases the credibility and transparency of the empirical design. I estimate the following equation set by two-stage least-squares:

$$y_{ie} = \alpha_1 + \nu f\left(r_{ie}\right) + \tau D_{ie} + \delta ticksize_{ie} + \varepsilon_{ie} \tag{8}$$

$$D_{ie} = \alpha_0 + \phi f\left(r_{ie}\right) + \psi T_{ie} + \varphi ticksize_{ie} + \varpi_{ie}, \tag{9}$$

where y_{ie} is the outcome of interest (e.g. liquidity, trading activity); $f\left(r_{ie}\right)$ is a global polynomial function of the ranking variable r_{ie}; D_{ie} is an indicator for anonymous trading; and $T_{ie} = \mathbf{1}\left[r_{ie} \geq 0\right]$, where r_{ie} has been normalized to zero by subtracting r_{ie} from 25 (see Section 4 for details). Notice that I do not include interaction terms between $f\left(r_{ie}\right)$ and D_{ie} or T_{ie}. Such interaction terms allow for a more flexible relationship between r_{ie} and y_{ie}, which may reduce the potential for bias in the RD estimates, but at the same time create an expanding set of endogenous variables that need to be instrumented. In this section, I sacrifice some flexibility in order to preserve statistical power. In Table A.3, I present estimates of τ from five models with different polynomial specifications of the relationship between r_{ie} and y_{ie}.

Table A.3: Robustness: Polynomials in r_{ie}

	Polynomial specification				
	Linear	*Quadratic*	*Cubic*	Quartic	Quintic
Dep. variable: Relative spread (log)					
τ	-0.53***	-0.54***	-0.64***	-0.64***	-0.64***
	(-3.84)	(-4.56)	(-3.83)	(-4.96)	(-3.17)
%Δ	-41.34	-41.70	-47.30	-47.06	-47.32
N	200	200	200	200	200
Dep. variable: Trades (log)					
τ	0.55**	0.56***	0.62**	0.61***	0.87***
	(2.47)	(3.22)	(2.19)	(2.85)	(2.87)
%Δ	73.24	75.30	86.16	84.57	138.92
N	200	200	200	200	200
Dep. variable: Trading volume (log)					
τ	0.44*	0.45*	0.65*	0.63***	0.89**
	(1.80)	(1.98)	(1.89)	(2.75)	(2.43)
%Δ	55.17	56.16	91.06	88.69	143.67
N	200	200	200	200	200
Dep. variable: Volatility					
τ	0.00	0.00	0.00	0.00	-0.00
	(1.17)	(1.17)	(0.01)	(0.01)	(-0.00)
%Δ	0.06	0.06	0.00	0.00	-0.00
N	200	200	200	200	200

Note: The table provides estimates of τ from five separate separate fuzzy regression discontinuity designs. The specification is estimated using the trader anonymity events in the period $2008 - 2010$. The second stage regression specification is $y_{ie} = \alpha + \nu f(r_{ie}) + \tau D_{ie} + \delta ticksize_{ie} + \varepsilon_{ie}$. τ is estimated in a two-stage least-squares (2SLS) approach, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. $T_{ie} = \mathbf{1}[r_{ie} \geq 0]$, where r_{ie} is a ranking variable determined each June and December in $2008 - 2010$, by previous six months trading volume. Exogenous controls include the ranking variable r_{ie} and $ticksize_{ie}$. The five models in this table are estimated using different polynomial specifications on the relationship between r_{ie} and outcomes y_{ie}, ranging from a 1st order polynomial (linear) to a fifth order polynomial (quintic). The 2SLS is estimated within a bandwidth $h = 25$. h indicates the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). Percentage change, %Δ, is calculated as $e^{\tau}-1$. Standard errors are clustered at r_{ie}. t-statistics in parentheses.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

A.4 Robustness test: Dynamic RD design

In the main analysis, I investigate how trader anonymity at the Oslo Stock Exchange (OSE) affects stock-level outcomes by using a regression discontinuity (RD) design. With the RD design, I find that trader anonymity increases stock liquidity (smaller relative bid-ask spreads) and trading activity (number of trades and traded value) but has no effect on returns volatility. The RD design I employ in the main analysis is static in the sense that it does not take into account that uptake of anonymous trading in one period potentially affects the probability of receiving anonymous trading in subsequent periods. Such dynamics can arise because 1) anonymous trading is assigned based on trading volume, and 2) anonymous trading increases trading trading volume (see Table 4).

In this appendix, I allow for such dynamic effects by estimating a dynamic fuzzy regression discontinuity design. The specification I employ is inspired by Cellini et al. (2010) and Cuñat et al. (2012) but takes into account that there is imperfect compliance to the trader anonymity assignment rule. I estimate the following equation set by 2SLS:

$$y_{ie} = \alpha_1 + f\left(r_{ie}\right) + \tau D_{ie} + \delta ticksize_{ie} + \left[\sum_{t=1}^{t=j}\left(\theta_{e-t} D_{i,e-t} + f\left(r_{i,e-t}\right)\right)\right] + \varepsilon_{ie} \qquad (10)$$

$$D_{ie} = \alpha_0 + f\left(r_{ie}\right) + \psi T_{ie} + \varphi ticksize_{ie} + \left[\sum_{t=1}^{t=j}\left(\Omega_{e-t} T_{i,e-t} + f\left(r_{i,e-t}\right)\right)\right] + \varpi_{ie}, \qquad (11)$$

where y_{ie} is some outcome (e.g. stock liquidity); stock i during event e is predicted for anonymous trading if $r_{ie} \geq 0$; and D_{ie} is an indicator variable for anonymous trading. Since there is imperfect compliance to the main assignment rule $r_{ie} \geq 0$, I use $T_{ie} = \mathbf{1}\left[r_{ie} \geq 0\right]$ as an instrumental variable (IV) for actual treatment D_{ie}. I include a full set of lags $D_{i,e-t}$, that are instrumented by the corresponding $T_{i,e-t}$, to account for the potential impact of previous treatment status on contemperaneous outcomes. Both D_{ie} and T_{ie} are constrained to zero for all events e before trader anonymity was introduced. I include a full set of lags in $f\left(r_{ie}\right)$ as exogenous regressors. Recall that r_{ie} was determined also in periods before trader anonymity was introduced, and its inclusion as a dynamic regressor controls for the impact of previous high or low rankings on current outcomes. $ticksize_{ie}$ is added to control for stock-level differences in tick size. The treatment effect τ in equation 10 can now be interpreted as the contemporaneous effect of anonymous trading in event e, net of effects operating through successive trader anonymity assignments.

I follow Cellini et al. (2010) and Cuñat et al. (2012) and estimate the dynamic RD design parametrically. To do so, I employ the polynomial expansion approach described in Appendix A.3 with a fifth order polynomial in $f(r_{ie})$. For transparency and robustness, I estimate the dynamic RD design separately for one, two, and three lags in D_{ie}, T_{ie}, and r_{ie}. Estimates of τ are presented in Table A.4. Notice from the table that the number of observations decreases in the number of lags applied. This is because more lags require a stock to have been eligible for trader anonymity, by being an OSEBX index listed stock, for consecutive periods. Consequently, the number of observations will be lower in the dynamic specification than in the baseline polynomal approach (Appendix A.3).

Table A.4: Robustness: Dynamic RD design

	Dynamic specification		
	One lag	*Two lags*	*Three lags*
Dep. variable: Relative spread (log)			
τ	-0.58***	-0.48*	-0.48*
	(-2.94)	(-1.81)	(-1.82)
%Δ	-44.25	-37.87	-38.31
N	189	184	177
Dep. variable: Trades (log)			
τ	0.76**	0.90**	0.96***
	(2.40)	(2.47)	(2.78)
%Δ	112.84	146.99	161.98
N	189	184	177
Dep. variable: Trading volume (log)			
τ	0.88**	1.02**	1.06**
	(2.35)	(2.49)	(2.63)
%Δ	142.25	178.48	187.50
N	189	184	177
Dep. variable: Volatility			
τ	-0.00	-0.00	-0.00
	(-0.36)	(-0.06)	(-0.20)
%Δ	-	-	-
N	189	184	177

Note: The table provides estimates of τ from a dynamic fuzzy regression discontinuity design. The specification is estimated using the trader anonymity events in the period $2008-2010$. The second-stage regression is $y_{ie} = \alpha_1 + f(r_{ie}) + \tau D_{ie} + \delta ticksize_{ie} + \left[\sum_{t=1}^{t=j} \left(\theta_{e-t} D_{i,e-t} + f(r_{i,e-t})\right)\right] + \varepsilon_{ie}$. τ is estimated in a two-stage least-squares (2SLS) approach, where predicted treatment status T_{ie} for stock i during event e, is used as an instrumental variable for actual treatment status, D_{ie}. Similarly, $T_{i,e-t}$ is used as an instrument for $D_{i,e-t}$. $T_{ie} = \mathbf{1}[r_{ie} \geq 0]$, where r_{ie} is a ranking variable determined each June and December in $2002-2010$, by previous six months trading volume. Exogenous controls include a full set of fifth order polynomials in r_{ie} and $ticksize_{ie}$. The 2SLS is estimated within a bandwidth $h = 25$. h indicates the number of stocks included on either side of the treatment cutoff ($r_{ie} = 0$). Percentage change, %Δ, is calculated as $e^{\tau}-1$. Standard errors are clustered at r_{ie}. t-statistics in parentheses.
** p < 0.10, ** p < 0.05, *** p < 0.01*

A.5 Generating $Frag_{ie}$

In Section 6.3, I use a variable $Frag_{ie}$ to control for the effect of order flow fragmentation on stock outcomes (e.g. stock liquidity) in the fuzzy regression discontinuity design. In this section, I describe how $Frag_{ie}$ is generated.

To generate $Frag_{ie}$, I use weekly frequency data on pan-European trading activity in all stocks at the Oslo Stock Exchange (OSE) in the period $2008-2010$. The data is obtained from Fidessa, a commercial provider of software, trading systems, and market data to both buy-side and sell-side investors. The Fidessa data provide weekly accounts of total trading volume, in both currency values and shares traded, and number of transactions, for all OSE stocks, separately for trading on all European trading venues. All trading venues in the data are defined by Fidessa as either lit order books (LIT), dark order books (DARK), over-the-counter (OTC), or systematic internaliser (SI). For unknown reasons, six OSE firms are missing from the Fidessa data. Their stock tickers are GOGL, SNI, STXEUR, AWO, WWI, and SAS NOK. Results are insensitive to treating these observations as missing, or as zeros. In Section 6.3, I treat these observations as missing; hence the smaller number of observations.

I capture order flow fragmentation by the share of trading volume that takes place on other trading venues than OSE. I make no distinction between trading on LIT, DARK, OTC, or SI trading venues. I define $Frag_{it}$ for stock i on date t as the share of currency trading volume that occurs on all trading venues excluding OSE, relative to the total currency trading volume across all trading venues including OSE. That is, if A^+ is the set of all trading venues v in the Fidessa data, including OSE, and A^- is the set of trading venues excluding OSE, then $Frag_{it}$ is defined as:

$$Frag_{it} = \frac{\Sigma_{v \in A^-} Volume_{it}}{\Sigma_{v \in A^+} Volume_{it}}. \tag{12}$$

I average this measure within each event e (defined in Section 3), and obtain $Frag_{ie}$. The results in Section 6.3 remain quantitatively similar if the sets A^- and A^+ include only LIT trading venues, or if the sets include both LIT and DARK trading venues. The results in Section 6.3 also remain similar if fragmentation is instead measured by a so-called Herfindalh-Hirschman Index (HHI).

A.6 Difference-in-differences

In the main analysis, I investigate how trader anonymity at the Oslo Stock Exchange (OSE) affects market quality by using a regression discontinuity (RD) design. With the RD design, I find that trader anonymity increases liquidity (smaller relative bid-ask spreads) and trading activity (number of trades and traded value) but has no effect on returns volatility. These results may be driven by other market structure developments than anonymity, taking place at the same time. For example, the main sample period ($2008 - 2010$) is characterized by, among other things, an explosion in high-frequency trading (e.g. Jørgensen et al. 2014, Angel et al. 2011,2013), aggressive use (by stock exchanges) of new fee structures, such as maker-taker fees (e.g. Malinova and Park 2015), and a financial crisis. If these developments systematically correlate with OBX list membership, they may bias my estimates.

To minimize the potential for time-varying confounders, I employ a short-run difference-in-differences specification, surrounding only the first assignment of anonymous trading, on June 2, 2008. On this date, anonymous trading was introduced for the 25 stocks in the OBX list, while trading in all other stocks remained non-anonymous. Shortly after, on June 20, the composition of anonymously and non-anonymously traded stocks was revised as part of a routine revision of the OBX list. Therefore, I exclude the period June 2 to June 17 and consider June 20, 2008 to be the 'treatment date' of interest. I estimate the following DiD specification surrounding this date:

$$Y_{it} = a + \nu D_t^{Post} + \gamma D_i^{Treatment} + \tau D_{it}^{Post*Treatment} + \delta ticksize_{it} + \varepsilon_{it}, \tag{13}$$

where $D_t^{Post} = 1$ for all dates t after June 20, 2008 and 0 otherwise. $D_i^{Treatment} = 1$ for the treatment group and 0 for the control group. I define the treatment group as the sample of stocks traded anonymously as of June 20, 2008, and the control group as the sample of stocks traded non-anonymously as of June 20, 2008. $D_{it}^{Post*Treatment}$ is the interaction between D_{it}^{Post} and $D_{it}^{Treatment}$ which equals 1 for anonymously traded stocks in the treatment period and 0 otherwise. I control for stock-level differences in tick size by including $ticksize_{it}$. The treatment effect of anonymous trading is given by the coefficient τ in equation 13.

An added benefit of this simple difference-in-differences approach is that it allows for direct comparability with previous empirical work on trader anonymity, where equation 13 is the preferred specification (e.g. Friederich and Payne 2014, Dennis and Sandås 2015).

For further comparability with this existing literature, I define my sample period similar to that used by Dennis and Sandås (2015). Particularly, I estimate equation 13 using three months of data before and after June 20, 2008. Friederich and Payne 2014, in contrast, employ a sample with six months of data before and after their anonymity introduction date. Estimating equation 13 using six months of data before and after June 20, 2008, instead of three months before and after this date, delivers similar coefficient estimates of τ.

I estimate the difference-in-differences model separately for bandwidths $h = 5$, $h = 10$, $h = 15$, $h = 20$, and $h = 25$. h indicates the number of stocks included on either side of the marginal OBX stock ($r_{ie} = 0$). For example, when $h = 5$, the sample is restricted to the 10 stocks closest to the marginal anonymously traded stock. Restricting the sample this way offers two benefits. First, it provides a homogeneous sample of stocks, where it may be plausible that the common trend assumption of the DiD specification is satisfied. Second, it offers transparency and robustness to the estimation. The drawback of this approach, of course, is that specifications with small h have few observations, which may lead to noisy estimates of τ. For this reason, I consider $h = 25$ to be the main sample.

Volatility, in previous sections defined as the variance of close-to-close returns, is now proxied by the daily high price divided by the daily low price in order to have variation on a daily frequency. The remaining outcome variables — relative bid-ask spreads, number of trades, and trading volume — are defined as in previous analysis but now on a daily frequency.

Table A.5: Robustness: Difference-in-differences

	Bandwidth				
	h=5	*h=10*	*h=15*	*h=20*	*h=25*
Dep. variable: Relative spread (log)					
τ	-0.31**	-0.11	-0.16*	-0.21***	-0.21***
	(-2.35)	(-1.07)	(-1.88)	(-2.86)	(-3.10)
%Δ	-26.71	-10.82	-14.65	-18.85	-19.08
Adj. R^2	0.17	0.22	0.27	0.33	0.38
N	1128	2258	3386	4516	5640
Dep. variable: Trades (log)					
τ	0.36	0.20	0.27	0.24*	0.34***
	(0.83)	(0.82)	(1.55)	(1.75)	(2.73)
%Δ	44.04	21.80	30.70	27.64	40.18
Adj. R^2	0.15	0.36	0.42	0.54	0.60
N	1129	2259	3387	4517	5641
Dep. variable: Trading volume (log)					
τ	0.07	0.12	0.27**	0.24*	0.24**
	(0.26)	(0.80)	(2.20)	(1.97)	(2.25)
%Δ	6.83	12.71	31.11	26.57	27.20
Adj. R^2	0.13	0.37	0.45	0.55	0.59
N	1129	2259	3387	4517	5641
Dep. variable: Volatility					
τ	0.01	-0.00	-0.00	-0.00	-0.00
	(0.58)	(-0.18)	(-0.52)	(-0.86)	(-0.42)
%Δ	-	-	-	-	-
Adj. R^2	0.02	0.05	0.06	0.07	0.06
N	1129	2259	3387	4517	5641

Note: The table provides estimates from a difference-in-difference specification surrounding the first assignment of trader anonymity at the Oslo Stock Exchange, on June 2, 2008. Due to a change in the composition of anonymously traded stocks on June 18, 2008, I exlude all dates between June 2 and June 18. The sample period is defined as March 20, 2008, to September 20, 2008. The regression specification is $Y_{it} = a + \nu D_{it}^{Post} + \gamma D_{it}^{Treatment} + \tau D_{it}^{Post*Treatment} + \delta ticksize_{it} + \varepsilon_{it}$. $D_{it}^{post} = 1$ for all time periods after June 18, 2008, 0 otherwise. $D_{it}^{treatment} = 1$ for stocks traded anonymously as of June 18, 2008, 0 otherwise. $D_{it}^{post*treatment}$ is the interaction between D_{it}^{Post} and $D_{it}^{Treatment}$, and equals 1 for anonymously traded stocks in the post-treatment period, and 0 otherwise. The difference-in-differences model is estimated separately for bandwidths $h = 5$, $h = 10$, $h = 15$, $h = 20$, and $h = 25$. h indicates the number of stocks included on either side of the the marginal OBX stock ($r_{ie} = 0$). τ gives the treatment effect of trader anonymity. %Δ gives the percentage treatment effect for log coefficiencts, $e^{\tau} - 1$. Standard errors are clustered at the stock-level.

** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

A.7 Trader classification

In Section 7, I explore the impact of trader anonymity on the trading activity of institutional and retail investors. In this appendix, I describe how traders are classified as institutional and retail.

The starting point of this trader classification is the transaction-level data described in Section 3. Each transaction in the data reveals the identity of both the buyer and the seller. Unlike some particularly detailed datasets — for example Barber et al. (2009) and Malinova and Park (2015) — the buyer and seller identities in my data are at the brokerage firm level and do not identify underlying accounts. This means that all inference on trader type will be based on observable characteristics at the brokerage level. van Kervel and Menkveld (2015) use a similar approach to identify high-frequency traders at the Nasdaq OMX.

The first step in the trader classification is to compile a list of brokerage firms that execute atleast one transaction during the sample periods defined in Section 3. Brokerages are identified in the data by ticker codes (e.g., XYZ). I translate all ticker codes into full brokerage firm names using membership lists obtained from the Oslo Stock Exchange (OSE). The final list holds 66 unique brokerage firms.

I proceed to hand-collect information on each active brokerage from company home pages, member descriptions at the OSE, and from various financial web sites such as Bloomberg Business. From these sources, I am able to infer, albeit noisily, whether a brokerage firm represents, for example, an investment bank catering to institutional or high-net-worth clients, such as Goldman Sachs or Deutche Bank, a market-maker, such as Knight Capital, or an online discount brokerage such as E-Trade.

I use this information to decompose the overall order flow into components of retail and institutional order flows. I begin by isolating order flow from online discount brokerages, who cater to individual investors. In total, I identify five active discount brokerages at the Oslo Stock Exchange in the period 2008 – 2010. These brokerages are Avanza Bank AB, E*Trade Danmark A/S, Net Fonds ASA, Nordnet Bank AB, and Skandiabanken AB.

The residual order flow, which, judging by brokerage firms' self-descriptions and Oslo Stock Exchange member descriptions, consists predominantly of investment banks catering to institutional clients, market makers, and high-frequency trading firms, is collectively referred to as 'institutional.' I follow Linnainmaa and Saar (2012) and further decompose the institutional order flow into components of domestic and foreign institutional order flows. Domestic brokerages include all Scandinavian brokerages or any foreign subsidiary registered as a Scandinavian company (*Aksjeselskap* (AS) *or Aktiebolag* (AB)). Brokerages head-quartered outside Scandinavia are considered foreign.

Chapter 2

Tick Size Competition, High Frequency Trading, and Market Quality

Tick Size Competition, High Frequency Trading, and Market Quality

Tom Grimstvedt Meling and Bernt Arne Ødegaard*

May 2017

Abstract

We show that competitive stock exchanges undercut other exchanges' tick sizes to gain market share, and that this tick size competition increases investors' trading costs. Our empirical analysis is focused on an event in 2009 where three stock exchanges, Chi-X, Turquoise, BATS Europe, reduced their tick sizes for stocks with an Oslo Stock Exchange (OSE) primary listing. We find that the tick size-reducing exchanges captured market shares from the large-tick OSE. Trading costs at the OSE increased while trading costs in the competing exchanges remained unchanged. High frequency trading appears to be the main driver behind the market share and trading cost results. Our findings suggest that unregulated stock markets can produce tick sizes that are excessively small.

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

*Meling is at the University of Bergen. Tom.Meling@econ.uib.no. Ødegaard is at the University of Stavanger. bernt.a.odegaard@uis.no. This paper has benefited from discussions with Tamás László Bátyi, Terrence Hendershott, Hans K. Hvide, Teis Lunde Lømo, Christine Parlour, Bjørn Sandvik, Eirik A. Strømland, and Jonas Tungodden. Discussants Gernot Doppelhofer, Carole Gresse, and Stig Lundeby have provided invaluable insights.

Introduction

Over the past two decades, regulatory reforms in the United States and Europe have facilitated increased competition between stock exchanges.[1] Competition between stock exchanges can benefit market participants by promoting more efficient trading services. However, competition can also harm market participants if there are negative externalities. This paper studies a situation where competition induces exchanges to implement market design changes that worsen trading conditions for market participants. Our empirical setting involves European stock exchanges and their choice of tick size — the smallest price increment on the exchange. We show that competitive stock exchanges undercut each other's tick sizes to gain market share, and that market participants' trading costs increase as a consequence.

How large should tick sizes be? The early theoretical literature concluded that the optimal tick size is small but not zero (e.g., Cordella and Foucault 1999; Foucault, Kadan, and Kandel 2005). A larger tick size increases the cost of undercutting the limit orders of other investors, which can give incentives for investors to provide liquidity with limit orders. Moreover, a larger tick size can force the quoted bid-ask spread to be artificially wide, providing incentives for traders to make markets and thus increase liquidity. Meanwhile, this increase in the minimum bid-ask spread also increases investors' trading costs, partly offsetting the liquidity gains from incentivizing market making. Hence, the optimal tick size involves a trade-off between increasing investors' trading costs and providing incentives for liquidity provision.[2]

Opening for competition between stock exchanges can put downward pressure on tick sizes. Buti, Consonni, Rindi, Wen, and Werner (2015) show theoretically that exchanges with small tick sizes can capture market shares from large-tick exchanges — potentially giving an incentive for competitive exchanges to undercut the tick sizes of other exchanges to gain market share. However, the exchanges in the Buti et al. (2015) model are restricted from strategically adjusting their tick sizes. For this reason, the model does not provide clear predictions about what tick size would arise endogenously through competition between stock exchanges, and whether the competitive tick size would increase or decrease market quality compared to the tick size in a non-competitive stock market. Absent theoretical predictions, empirical work may provide guidance about the mechanisms through which competition can affect exchanges' tick size choice and market quality.[3]

[1] In the United States, the Regulation National Market System (Reg NMS) was introduced in 2005, while the Markets in Financial Instruments Directive (MiFID) was implemented in Europe in late 2007. Both Reg NMS and MiFID introduced new rules that intensified competition between trading platforms. For example, MiFID opened for competition between European stock exchanges by abolishing the so-called 'concentration rule', which previously forced all regulated trades to be executed in specific domestic marketplaces.

[2] The tick size is currently among the most controversial market design features in the current equity market policy debate, as market regulators in the United States and Europe are considering comprehensive market design reforms in search of a suitable tick size. For example, market regulators in the U.S. have recently implemented a large-scale pilot program that will increase the tick size for 1200 randomly chosen securities. The current proposal by European regulators is that tick sizes should be stock-specific, and be determined as a function of both the stock price and the stock liquidity.

[3] Tick sizes are heavily regulated in many of the world's most important stock markets, which may partly explain why the existing theoretical literature has yet to explore the consequences of having market forces determining the tick size. For example, the U.S. market regulator mandates a fixed tick size at \$0.01 for most

The purpose of this paper is to empirically assess the impact of opening for competition on exchanges' choice of tick size, and the consequence of competitive tick size choices for market quality. To this end, we study exchanges' strategic tick size decisions for Oslo Stock Exchange (OSE) listings in the aftermath of the MiFID reform, which in November 2007 opened for competition between European stock exchanges.[4] We focus on an event where three entrant exchanges, Chi-X, Turquoise, and BATS Europe reduced the tick size for their selections of OSE listed stocks. Chi-X moved first, and reduced the tick size on June 1, 2009. Turquoise and BATS quickly followed and reduced their tick sizes on June 8 and June 15, respectively. The OSE responded within a month by reducing its own tick size. This race to the bottom ended when the Federation of European Securities Exchanges (FESE) brokered a common tick size across all the exchanges, mandating much smaller tick sizes than before the 'tick size war.'

We leverage extremely rich data on the trading of OSE listed stocks across all European trading platforms to explore why opening for competition between exchanges seems to drive tick sizes down. Our findings suggest that reducing the tick size can be an effective strategy for entrant exchanges to increase their market share. In particular, we find that Chi-X nearly doubled its market share of overall trading from the first day with reduced tick sizes. In contrast, the late-movers Turquoise and BATS Europe were unable to capture market shares from the OSE with similar tick size reductions. Likewise, when the OSE retaliated and reduced its own tick sizes, it was unable to reclaim the lost market share. Thus, our findings suggest that competitive stock exchanges have a strong incentive to undercut other exchanges' tick sizes, as such tick size competition can permanently increase their market share.

Our data also allow us to estimate the impact of tick size competition on measures of market quality at individual trading platforms. Using a difference-in-differences approach, we find that tick size competition negatively affected stock liquidity at the OSE and Chi-X — the two exchanges with market share gains or losses during the tick size war in June 2009. Our empirical strategy is to compare changes in stock liquidity for stocks that were directly affected by the tick size war (stocks listed at both the OSE and Chi-X) to changes in stock liquidity for stocks unaffected by the tick size war (stocks listed only at the OSE). We find that trading costs at the OSE increased after the Chi-X tick size reduction, while trading costs at Chi-X remained unchanged, suggesting an overall increase in trading costs. We also find that order book depth at both Chi-X and the OSE suffered greatly from the OSE retaliatory tick size reduction. Our results persist after controlling for stock-level changes in trading volume, suggesting that the observed changes to stock liquidity cannot fully be explained by a redistribution of trading volume between exchanges.

---

securities and stock exchanges. In Europe, the proposed MiFID II legislation will enforce a common tick size regime across exchanges that compete for the same order flow. The pervasiveness of tick size regulations in stock markets around the world also means there are few empirical settings that researchers can analyze to understand the strategic tick size choices of competitive stock exchanges.

[4]Before the implementation of MiFID in 2007, the OSE was the monopolist marketplace for the trading in stocks with an OSE primary listing. After the MiFID reform, new exchanges quickly entered to offer trading in OSE stocks. These entrant exchanges long struggled to get a toe-hold in the market, but competition had slowly taken hold by early 2009. Section 1 provides further details on the MiFID reform, the ensuing increase in competition for OSE listed stocks, and tick size regulations in Europe.

To explore the mechanisms through which small-tick exchanges capture market share, and to assess why tick size competition seems to decrease market quality, we leverage very detailed order book data from the OSE. A key theoretical result in the Buti et al. (2015) model is that traders migrate to small-tick markets because the bid-ask spread is constrained by the tick size in the large-tick market. The mechanism behind this result is that a constrained bid-ask spread makes it harder for traders to undercut limit orders to gain execution priority, which induces impatient traders to send their orders to an exchange where the tick size is smaller and undercutting is easier. Inconsistent with this theoretical prediction, we find that the extent of OSE market share loss during the tick size war is unrelated to the severity of bid-ask spread constraints at the OSE. In fact, few stocks in our sample trade with bid-ask spreads that are close to being constrained by the tick size.

Rather than constraints to bid-ask spreads, preliminary results suggest that high-frequency traders (HFTs) appear to be responsible for the observed redistribution of market share from large-tick to small-tick exchanges. We generate a stock-level proxy for HFT activity (the 'order to trade ratio'), and find that OSE stocks with more HFT activity experienced a greater loss in market share following the Chi-X tick size reduction. To further investigate this mechanism, we show that certain traders, who we conjecture are HFTs, migrated the OSE in favor of Chi-X to execute at prices that were unattainable with the coarse price grid at the OSE. Finally, we find that HFTs became much more active at the OSE after the OSE tick size reduction, illustrating that HFTs prefer to trade when tick sizes are small.

We offer a tentative mechanism through which HFT order flows can account for the observed changes to stock liquidity.[5] Since stock liquidity at Chi-X seemingly did not improve from an inflow of HFT volume, we conclude that the HFTs that migrated to Chi-X traded as liquidity-demanders or alternatively that these HFTs were informed investors whose trading imposed an adverse selection cost on limit order traders at Chi-X. However, given the observed increase in trading costs at the OSE in the same period, the same HFTs appear to improve liquidity when they trade at the OSE. We interpret this finding as consistent with HFTs switching from trading as liquidity-providers in the large-tick OSE market to trading as liquidity-demanders in the small tick Chi-X market.

Our paper contributes to several threads in the current academic debate over optimal tick sizes in equity markets.[6] First, a recent empirical literature studies how a regulatory-mandated tick size difference between over-the-counter markets ('dark pools') and regular exchanges in the United States affects the order-routing decisions of investors (e.g. Bartlett and McCrary 2015, Kwan, Masulis, and McInish 2015, Buti et al. 2015). Consistent with this literature, we find that investors send their orders to trading platforms that allow for trading at smaller tick sizes. However, we add to this literature by exploring the tick size that arises endogenously through competition between exchanges that can strategically adjust their tick size, and estimate the

[5] A future version of this paper will test this mechanism more formally.

[6] For recent surveys of the voluminous empirical and theoretical academic literatures on the role of tick sizes in equity markets, see Holden, Jacobsen, and Subrahmanyam (2013), Securities and Exchange Commission (2012) and Verousis, Perotti, and Sermpinis (2017).

effects of this competitive tick size on market quality.

Second, our findings seem to contradict the empirical literature which shows that HFTs trade more actively when tick sizes are large. For example, O'Hara, Saar, and Zho (2015) and Yao and Ye (2015) provide empirical evidence that HFTs are more active in liquidity provision and have larger profit margins when tick sizes are large. The mechanism that the authors propose is that the HFT speed advantage becomes more valuable when price competition is constrained by the tick size. Our results, in contrast, suggest that HFT seem to migrate the large-tick OSE in favor of small-tick competing exchanges, indicating an opposite HFT preference over tick sizes. These conflicting results may suggest that certain types of HFT strategies may require a small tick size whereas other HFT strategies, such as liquidity-provision, may require a larger tick size.

Finally, our results provide empirical support for the current market regulations in the United States that enforce a common tick size across competing exchanges, and for the proposed regulations in Europe that aim to accomplish the same (see footnote 3). Specifically, our results show that individual stock exchanges have an incentive to reduce their tick sizes to capture market shares and, at the same time, that such tick size reductions can have negative effects on the stock liquidity in competing marketplaces. Thus, a conceivable consequence of tick size competition is that combined market liquidity (across all trading venues) declines. Market regulators can restrict stock exchanges' ability to engage in destructive tick size competition by enforcing a common tick size regime across all exchanges competing for the same order flow.

The paper proceeds as follows. Section 1 provides institutional background on equity trading at the Oslo Stock Exchange and describes the tick size war for OSE listed stocks; Section 2 develops testable theoretical hypotheses; Section 3 describes our data; Section 4 studies the impact of tick size competition on the distribution of market shares across exchanges; Section 5 estimates the impact of tick size competition on market quality; Section 6 explores the mechanisms that link tick size competition to market fragmentation and market quality; and Section 7 concludes.

1 Institutional Background

This paper explores exchanges' strategic tick size decisions for Oslo Stock Exchange listings in the aftermath of the MiFID reform, which in November 2007 opened for competition between European exchanges. We focus on a series of tick size reductions for OSE listed stocks during the Summer of 2009, which we collectively refer to as the 'tick size war.' In this section, we first provide institutional details concerning the trading in Norwegian equities — both at the Oslo Stock Exchange and at competing trading platforms — before we summarize the events of the 'tick size war' in 2009.

1.1 The Oslo Stock Exchange

The Oslo Stock Exchange is a medium-sized stock exchange by European standards, currently ranking among the 30 largest (by market capitalization) equity markets in the world. At the end of 2010, the combined market capitalization of the OSE was about 1.8 trillion NOK, distributed across 239 companies. Over the last decade, the OSE has collaborated and shared trading technology with other European stock exchanges.[7] The collaboration with other exchanges has implied the use of common technology and, to some extent, common market models. Nevertheless, the OSE has remained relatively free to implement individual trading rules and compose an individual market model.

The OSE operates a fully computerized limit order book, and has done so since January 1999. The order book allows for conventional limit orders, market orders, iceberg orders and various other common order types. As is normal in electronic order-driven markets, order placements follow price-time priority — incoming orders are first sorted by their price and then, in case of equality, by the time of their arrival. The trading day at the OSE consists of three sessions: an opening call period, a continuous trading period, and a closing call period. Call auctions may also be initiated during continuous trading if triggered by price monitoring or to restart trading after a trading halt.[8]

The distributions of firm size and trading volume at the OSE are both heavily skewed. The OSE is dominated by a few very large companies. For example, the most valuable listed company, Statoil (an oil company), accounted in 2009 for about 25% of the OSE market capitalization. Two other companies, Telenor (telecommunications) and Den Norske Bank (integrated financial) each accounted for about 10% of the total market capitalization of the OSE. The large companies at the OSE also dominate in terms of trading activity. A considerable portion of overall trading volume takes place in the largest stocks at the OSE, and in particular in the constituent stocks of the large-cap OBX index. The OBX index comprises at any point of time the 25 most-traded (and typically the most valuable) stocks at the OSE.[9]

1.2 Competition for European order flow (MiFID)

Competition for European order flow is a fairly recent phenomenon. National stock exchanges, such as the Oslo Stock Exchange, long operated as monopolist marketplaces for trading in domestic shares. However, the introduction in 2007 November of the Markets in Financial Instruments Directive (MiFID) legislation unleashed competition for European order flow by abolishing the so-called 'concentration rule', which forced any regulated trade to be executed

[7] In 2002, the OSE introduced the SAXESS trading platform in cooperation with NASDAQ OMX. In 2009, the OSE partnered with the London Stock Exchange Group (LSEG) and implemented their TradElect trading platform in April 2010. The OSE now employs the Millennium trading system — the same trading system used by, for example, the London Stock Exchange and Borsa Italiana.

[8] For details on the trading fees and market transparency at the OSE, see for example Jørgensen, Skjeltorp, and Ødegaard (2016) or Meling (2016).

[9] The composition of the OBX index is revised twice a year, in June and December, primarily based on total stock trading volume at the OSE over the previous six months. Meling (2016) provides more details on the OBX index.

in the primary market. Today, European equity trading is scattered across a large number of trading venues that compete vigorously to attract order flow.

Three types of trading venues have emerged to compete for European order flow — Regulated Markets (RMs), Multilateral Trading Facilities (MTFs), and Systematic Internalisers (SIs). The RMs (such as the OSE) and the MTFs share similar features. For example, both RMs and MTFs can decide on the type of orders allowed on their order books, the structure of member fees (e.g. fixed, variable, maker-taker), and to some extent the transparency of the trading process. Moreover, both RMs and MTFs are allowed to organize primary listings. In practice, however, MTFs do not offer primary listing services, and can be viewed as the European equivalent of ECNs in the United States. Distinct from both RMs and MTFs, the SIs are investment firms that systematically match client orders internally or against their own accounts.

Some stylized facts based on publicly available data from Fidessa, a data vendor, may help understand MiFID's impact on the trading of OSE listed stocks. At the time of writing, in 2016, more than twenty regulated markets, multi-lateral trading facilities, systematic internalisers, or unregulated over-the-counter trading venues offer trading in the most liquid stocks at the Oslo Stock Exchange. The OSE retains the largest market share, followed by BATS over-the-counter (OTC), BATS CXE (formerly known as Chi-X), Turquoise, and BATS BXE (formerly known as BATS Europe). The OSE market share of overall trading (including over-the-counter trading) in its most liquid stocks has declined from 100% in 2007 to close to 40% in 2016.

1.3 OSE competitors: Chi-X, Turquoise and BATS

Three MTFs — Chi-X, Turquoise and BATS Europe — feature prominently in our study due to their proclivity to adapt their market designs to capture market shares. Established in 2007 by a consortium of investment banks, Chi-X was the first MTF in Europe. Both BATS Europe and Turquoise were established in 2008 — BATS by BATS Global Markets, a U.S. exchange operator, and Turquoise by a consortium of investment banks. In December, 2009, the London Stock Exchange Group acquired a 60% stake in the Turquoise platform. After our sample period, in 2011, BATS Europe has acquired Chi-X.

Similar to the OSE, Chi-X, Turquoise and BATS operate fully electronic matching engines where anonymous orders are matched continuously, according to price-time priority. Unlike the OSE, the MTFs aggressively employ maker-taker fees to incentivize liquidity supply. For example, at Chi-X, liquidity demander (takers) pay a transaction fee of 0.3 basis points while liquidity suppliers (makers) earn a rebate of 0.2 basis points.

Chi-X, Turquoise and BATS Europe offer trading in some, but not all, of the 200–300 stocks listed at the OSE. The three MTFs initially opened trading in only the largest and most liquid stocks at the OSE, before gradually expanding their selection. For example, Chi-X initially offered trading in only the five largest stocks at the OSE. By 2015, Chi-X offers trading in more than 50 OSE products. Similarly, Turquoise initially opened trading in 28 OSE stocks but has since greatly expanded its selection to by 2015 include more than 150 OSE products.

1.4 'Tick size war' for OSE listed stocks

The introduction of MiFID in November 2007 opened for competition between European trading platforms. However, the MiFID reform did not specify regulations concerning exchanges' choice of tick size — the smallest price increment on a stock exchange. This allowed competitive European exchange operators to strategically adjust their own tick sizes.[10] The purpose of our paper is to analyze an event where three entrant trading platforms, Chi-X, Turquoise, and BATS Europe unexpectedly in June 2009 decided to reduce the tick size for several of their stock listings.[11] The entrants' unexpected tick size reductions sparked a frenzy of tick size reductions which commentators at the time called a 'tick size war.'

The tick size war during the Summer of 2009 can conveniently be divided into three phases. In the first phase, which we call the *break-out phase*, Chi-X, Turquoise and BATS challenged the market positions of the Scandinavian primary markets (Oslo, Stockholm, and Copenhagen) by successively reducing the tick size for their selection of Danish, Norwegian, and Swedish stocks. The tick size war began on June 1, 2009, when Chi-X reduced its tick size. Turquoise followed on June 8, reducing the tick size for Scandinavian stocks as well as for five London listed stocks. Finally, BATS Europe reduced the tick sizes for Scandinavian stocks, ten London stocks, and five Milan stocks on June 15 (BATS, 2009).

The tick size reductions by Chi-X, Turquoise, and BATS during the *break-out phase* were substantial. In Table 1, we summarize the tick size schedules used by all four stock exchanges throughout the calendar year 2009. At the time of the Chi-X tick size reduction, on June 1, 2009, the OSE operated with three tick size schedules: a flat tick size of NOK 0.01 for Statoil (the most liquid stock at the OSE); a general tick size schedule for all OBX shares, with tick sizes varying between 0.01 and 0.25; and a separate tick size schedule for all illiquid (non-OBX) shares. The new Chi-X tick size schedule, in contrast, introduced a NOK 0.001 tick size for all OSE stocks traded at Chi-X with prices below NOK 10 and a NOK 0.005 tick size for stocks priced above NOK 10. The tick size schedules introduced by Turquoise and BATS were less aggressive, but they still offered substantially smaller tick sizes than the OSE.[12]

In the second phase of the tick size war — the *retaliation phase* — the OSE responded

[10] That European trading venues can determine their own tick sizes contrast with the regulatory setting in the United States. The U.S. market regulator (the Securities and Exchange Commission) mandates a fixed tick size for all stocks priced above $1 of $0.01.

[11] In the absence of formal tick size regulations after the MiFID reform, the Federation of European Securities Exchanges (FESE) brokered in March 2009 a 'gentlemen's agreement' between several European stock exchanges and MTFs to implement a common tick size regime. The motivation behind the tick size agreement was that individual trading venues can capture market shares by reducing their tick sizes but that such tick size competition can have a detrimental effect on stock liquidity (FESE 2009). The March 2009 tick size agreement involved four alternative tick size schedules that should determine a stock's tick size as functions of the stock price. However, the agreement did not clarify which of the four tick size schedules should be used, when the tick size schedules should be implemented, or who should make these decisions. Evidently, this ambiguous 'gentlemen's agreement' was insufficient to prevent Chi-X, Turquoise, and BATS from reducing their tick sizes.

[12] We can point out that prior to the tick size war, tick sizes for stocks listed at the OSE were large compared to the current penny tick size in the United States. For example, converted at the 2009 exchange rate of 6.3 NOK per USD, the pre-tick-size-war tick size of NOK 0.01 for Statoil translates into 0.15 cents. However, the post-war Chi-X tick size of 0.005 translates to only 0.08 cents. Thus, the tick size war pushed tick sizes for OSE listed stocks below the current US tick size regime.

in kind to its tick size reducing competitors. On July 6, 2009, the OSE reduced its tick size uniformly to NOK 0.01 for the 25 stocks in the OBX index. In a press release, the OSE declared that other trading 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." Doing so, the OSE largely mitigated the between-exchange tick size differences that arose during the *break-out phase.*

What can explain the exchanges' decisions to reduce their tick sizes during the Summer of 2009? First, to understand the strategic decision the OSE faced following its competitors' tick size reductions, we give a preview of our results concerning the OSE market share in its own stock listings. Figure 1 compares the distributions of daily market shares for the OSE and Chi-X before (May 2009) and after (June 2009) the Chi-X tick size reduction. The figure illustrates a sizable shift of market shares from the large-tick OSE market to the small-tick Chi-X market. More precisely, in Section 4.1 we estimate the OSE market share loss after the Chi-X tick size reduction to nearly three percentage points. Observing this rapid decline in market share, it is straight-forward to understand why the OSE found it 'necessary' to respond to competing exchanges' tick size reductions. Similarly, entrants may have an incentive to drive tick sizes further down, as this strategy seems to enable them to gain market share.

Second, contemporary observers argued that the exchanges' decisions to reduce their tick sizes were rooted in pressure from influential high-frequency trading (HFT) firms who desired smaller tick sizes (e.g., *Financial Times* 2009). As a preliminary exploration of this hypothesis, Figure 2 plots the order-to-trade ratio (OTR) separately for OBX index stocks at the OSE who were exposed to the July 6, 2009 OSE tick size reduction and non-OBX index stocks who were not exposed to the tick size reduction. The OTR is a commonly used proxy for HFT activity, and we define this proxy in more detail in Section 3.3. Consistent with HFTs wanting to trade in small-tick markets, Figure 2 shows a remarkable increase in HFT activity for OSE stocks affected by the July 6, 2009 tick size reduction.[13]

The final stage of the tick size war is the *harmonization* phase. On June 30, 2009, the FESE brokered a harmonization of tick sizes between the stock exchanges and the MTFs. FESE argued that the recent tick size reductions were not in the interest of end investors and that too granular prices could have detrimental effects on stock market depth. The FESE agreement facilitated a pan-European harmonization of tick size schedules for the most actively traded stocks, which significantly simplified and reduced the number of different tick size schedules used by the exchanges. The far right panel of Panel A in Table 1 displays the tick size schedule chosen by the OSE. These changes were to be implemented within two weeks and six months depending on the needs of the exchange. The Scandinavian markets responded in steps. OSE harmonized tick sizes August 31, 2009. The other markets followed later, Stockholm on October

[13]It is useful to point out the parallels between our analysis and Menkveld (2013), who explores the entry of a HFT market maker in the Dutch stock market in the beginning of 2008. Anecdotal evidence suggests that the market maker in Menkveld (2013), Getco, and other similar trading firms, gradually expanded their operations into other European marketplaces. The increase in HFT activity at the OSE in July 2009 can therefore indicate the entry of new HFTs in the Norwegian stock market.

26 and Copenhagen on January 4, 2010.[14]

2 Hypothesis development

Theoretical work in the equity market microstructure literature predicts that between-exchange differences tick size differences can influence investors' order-routing decisions and measures of stock liquidity. This section discusses the potential mechanisms through which the tick size war for OSE listed stocks (Section 1.4) can affect these stock market outcomes. To simplify the exposition, we assume the following sequence of mechanisms: First, there is an exogenous shock to the tick size at exchange v while the tick sizes at exchanges v^- remain unchanged. Second, investors reconsider whether to route their orders to exchanges v or v^-. Third, stock liquidity in each of the exchanges is affected directly by the choice to reduce the tick size (exchange v) and indirectly by investors' order-routing decisions (both exchanges v and v^-).[15]

Distribution of trading volume across exchanges: We present two mechanisms through which between-exchange tick size differences can affect affect investors' order-routing decisions, and subsequently alter the distribution of trading volume across stock exchanges. These mechanisms are motivated by two different strands of academic literature. First, recent theoretical work suggests that between-market tick size differences can shift trading volume from large-tick markets to small-tick markets. For example, Buti et al. (2015) predict that when a large-tick market faces competition from a small-tick market, some traders with access to both markets will route their orders to the small-tick market. The mechanism which generates their theoretical result is that large tick sizes make it more difficult for traders to undercut orders in the limit order book to gain execution priority. This induces impatient traders to route their orders to markets where price competition is less constrained by the tick size and undercutting is easier. A key prediction in Buti et al. (2015) is therefore that between-exchange tick size differences are more important for stocks where price competition is constrained by the tick size than for stocks where price competition is unconstrained.

The second mechanism we consider is that high-frequency traders (HFTs) and non-HFTs may react differently to changes in the tick size. For example O'Hara et al. (2015) and Yao and Ye (2015) argue that HFTs are more active in liquidity provision and have larger profit margins in a large-tick size environment. They argue that the HFT speed advantage becomes more valuable when price competition is more constrained by the tick size. By this logic, one should expect that HFTs react to the tick size reductions during the tick size war by routing their

[14] For a short while, the FESE tick size agreement successfully warded off competitive tick size reductions. However, in 2011, Euronext decided to implement a smaller tick size than agreed upon in the FESE agreement for certain liquid stocks, sparking "outrage" among competing trading platforms amid concerns of a new tick size war (e.g. *Financial Times* 2011). As a response to the seemingly unstable tick size agreements in Europe, the updated MiFID II regulation is expected to mandate a common tick size regime across all European trading platforms.

[15] We need to assume a sequence of mechanisms because, as econometricians, we only observe the initial shock to tick sizes and the simultaneous outcomes that correspond to step two (order-routing decisions) and step three (stock liquidity). This means that we cannot disentangle empirically whether tick size-induced changes to order-routing decisions causally affect stock liquidity, or whether tick size-induced changes to stock liquidity causally affect order-routing decisions.

orders to large-tick size exchanges instead of small-tick size exchanges, and thereby influence the distribution of market shares across exchanges.

However, other HFT strategies than liquidity-provision may become more profitable when tick sizes are small than when they are large. For example, cross-market arbitraging strategies rely on small and fleeting price discrepancies for the same security at different exchanges. A reduction in the tick size in one exchange means the increments by which prices can move will differ between exchanges, giving HFTs more opportunities to seek out trading opportunities across-exchanges. A different HFT strategy involves reacting to the arrival of new and valuable information before other traders have time to modify their previous (now mispriced) offers to buy or sell (Menkveld, 2016). This strategy may be easier to implement in small-tick markets as a reduction in the tick size lowers the marginal cost of undercutting existing quotes. In other words, we expect the extent to which HFTs prefer to route their orders to large-tick or small-tick markets to depend on the trading strategies that HFTs follow.

Stock liquidity in each of the exchanges: The tick size war for OSE listed stocks can also affect measures of stock liquidity in each of the involved stock exchanges. We conjecture that the overall impact of the tick size war on stock liquidity can be separated into three components. The first component is the same-market effect from reducing the tick size. Inspired by a voluminous empirical and theoretical literature on the impact of tick size reductions in monopolist limit order books, our baseline prediction is that stock exchanges that reduce their tick sizes should experience tighter bid-ask spreads and shallower order books (e.g., Securities and Exchange Commission 2012).

The second component of the overall effect of the tick size war on stock liquidity comes from the changing distribution of trading volume across exchanges. Exchanges that reduce their tick sizes may experience inflows of trading volume from exchanges that keep large tick sizes. Inflows (or outflows) of trading volume can improve or degrade stock liquidity, depending on the characteristics and trading strategies of the investors that migrate between exchanges. For example, reducing the tick size may cause wider (narrower) bid-ask spreads if it leads to an inflow of informed (uninformed) investors, on account of the greater (smaller) adverse selection costs faced by liquidity providers (e.g. Glosten and Milgrom 1985 or Kyle 1985). Similarly, if between-exchange tick size differences affect the order-routing decisions of HFTs, an inflow or outflow of HFT trading volume can improve or degrade stock liquidity, depending on whether the HFTs engage in market-making activities or conversely demand or degrade liquidity.[16]

The final theoretical mechanism we consider concerns the potential disruption of network externalities in liquidity provision, along the lines of Pagano (1989). Loosely speaking, a consolidated market that is already liquid can attract even more liquidity because of positive network

[16]Empirical evidence suggests that a majority of HF traders behave as market makers, with a business model of providing liquidity, compensated by the bid-ask spread, which can improve stock liquidity (e.g. Menkveld 2013 and Hagströmer and Nordén 2013). However, the empirical evidence also point to the presence of other forms of HFTs, who for example, use their speed advantage to "snipe" stale quotes before other traders can modify them. Another hypothesized HF strategy involves predicting future order flow, trying to determine the presence of large trades being worked over time, and trading in front of these. Some HFT strategies even resemble illegal price manipulation: for example the "spoofing" strategy involves filling the order book with orders away from the best bid and/or ask in order to manipulate other traders' order placement strategy.

externalities. This is because each additional trader in the liquid market reduces the search and trading costs for other potential traders, which attracts even more traders. Conversely, traders may be discouraged from entering an illiquid market because of high search and trading costs, which further degrades the illiquid market's liquidity (a negative network externality). The presence of such network externalities implies in our setting that an inflow (outflow) of trading volume at the liquid Oslo Stock Exchange can be relatively more beneficial (detrimental) to stock liquidity than a corresponding inflow or outflow of trading volume at the fairly illiquid MTFs.

Summary: This section discusses mechanisms through which the tick size war for OSE listed stocks can affect stock market outcomes. To summarize, we expect the tick size war to shift trading volume and market share from the large-tick size OSE exchange to its small-tick size competitors. This shift in market shares should be motivated by constraints to price competition at the OSE or by changes in the order-routing decisions of HFTs, or a combination of these two mechanisms. For the exchanges that reduce their tick size, we expect the direct effect to be narrower bid-ask spreads and shallower order books. This direct effect will be amplified or weakened by inflows of trading volume, depending on whether the migrating traders are informed or uninformed, and whether the migrating traders supply or consume liquidity. For the exchanges that maintain large tick sizes (the OSE), we expect that stock liquidity is affected through an outflow of trading volume and from the disruption of liquidity externalities. Sections 4 to 6 test these mechanisms empirically.

3 Data

This section presents the data we use to explore the impact of the tick size war between the Oslo Stock Exchange, Chi-X, Turquoise, and BATS, on the distribution of market shares across exchanges and the quality of trading in each of the exchanges. The section also defines our main outcome variables, and presents descriptive statistics of stock trading at the Oslo Stock Exchange, Chi-X, Turquoise and BATS.

3.1 Data Sources

We use several datasets in our empirical analysis. First, we use proprietary order-level data obtained from the 'market surveillance' group at the OSE. This dataset contains information on all orders submitted to the exchange, regardless of whether the order is executed or not. Orders are flagged indicating whether they are executed (a trade), canceled, or modified. The fact that we observe individual orders, not just the trades, allows us to calculate empirical measures of high-frequency trading activity, such as the "order-to-trade" ratio (equivalently, the "quote-to-trade" ratio).

Second, to analyze trading in OSE listed stocks on competing stock exchanges, we use the ThomsonReuters Tick History (TRTH) Database. The TRTH contains trade-and-quote data for OSE listed stocks across all European equity market places. For lit market places (markets with displayed order books) the dataset provides information on the ten best levels of the bid and ask side of the limit order book. The ThomsonReuters data also includes information on over-the-counter trading of OSE shares through the inclusion of trades reported by Markit BOAT (a MiFID-compliant trade reporting facility). We use the TRTH database to compute each stock exchange's market share of trading, as well as a wide range of stock liquidity measures (defined in Section 3.3).

Finally, we supplement these two datasets with information on end-of-day prices, OBX index constituency, and tick size levels, obtained from the Oslo Stock Exchange Information Service (OBI).

3.2 Sample restrictions

In our empirical analysis, we focus exclusively on stocks with a primary listing on the Oslo Stock Exchange (OSE) for which we have detailed data on the trading process. We restrict the sample period to the calendar year 2009, which encompasses all the relevant tick size changes (see Section 1.4). We restrict our attention to the trading that occurs on the OSE, Chi-X, Turquoise, and BATS Europe order books, as these were the four exchanges involved in the tick size war.

Throughout most of the empirical analysis, we restrict our sample to stocks in the large-cap index at the OSE, the OBX index. Only OBX index stocks were affected by the July 6, 2009 tick size reduction by the OSE. Moreover, though Chi-X, Turquoise, and BATS offered also offered trading in non-OBX stocks, most of their trading activity was focused on OBX index stocks. For this reason, our main sample comprises the 26 individual stocks in the OBX index.[17] We will in some of our analyses expand the sample to include all OSE listed stocks. This allows us to compare OSE listed stocks that were affected by the tick size changes to corresponding stocks unaffected by the tick size war.

3.3 Variable definitions

We explore the impact of the tick size war between the Oslo Stock Exchange, Chi-X, Turquoise, and BATS, on a number of common measures of stock market quality. To measure the transaction cost dimension of stock liquidity we use four spread measures of liquidity. First, the *relative spread* is defined as the difference between the current best bid and ask divided by the quote midpoint. We update the relative spread whenever the limit order book is updated, and calculate the average of these estimates throughout the trading day.

Second, the *effective spread* captures the cost of demanding liquidity. We define the effective proportional half-spread for trade j in stock i as $q_{ji}(p_{ji} - m_{ji})/m_{ji}$, where q_{ji} is an indicator variable that equals $+1$ for buyer-initiated trades and -1 for seller-initiated trades; p_{ji} is the trade price; and m_{ji} is the quote midpoint prevailing at the time of the trade. To determine whether an order is buyer or seller initiated, we compare the transaction price to the previous

[17] One stock (RCL) moves into the OBX index and another (AKER) moves out of the OBX index during the sample period (the relevant OBX revision date is June 19, 2009). We do not remove these stocks from the sample.

quote midpoint — if the price is above (below) the midpoint we classify it as a buy (sell). We compute average effective spreads across all transactions during the trading day.

Third, the *realized spreads* measure the gross revenue to liquidity suppliers after accounting for adverse price movements following a trade. The 5-minute realized spread for transaction j in stock i is given by $q_{ji}(p_{ji} - m_{i,j+5\text{min}})/m_{ji}$, where $m_{i,j+5\text{min}}$ is the quote midpoint 5 minutes after the j'th trade. q_{ji} and p_{ji} are defined as before. Similar to the effective spread, we calculate the daily average of realized spreads for all trades during the day.

Fourth, the *price impact* captures the gross losses to liquidity demanders due to adverse selection. The five-minute price impact for a given transaction j in stock i is defined as $q_{ji}\left(m_{i,j+5\text{min}} - m_{i,j}\right)/m_{ji}$. We calculate our measure of price impact at the stock-day level by averaging the price impact across all trades during the trading day.

We estimate the *depth* of the limit order book by calculating the sum of pending trading interest at the best bid and ask prices. Our measure of order book depth is updated whenever the limit order book is updated, and averaged across all order book states throughout the trading day. To proxy for the noise in the price process, we estimate *realized volatility* as the second (uncentered) sample moment of within-day ten-minute returns.

We use the so-called order-to-trade ratio (OTR) to proxy for the extent of high-frequency trading activity at the stock-day level. The OTR is the ratio of messages (orders, order cancellations, order modifications) submitted to the exchange's limit order book relative to the number of completed transactions. As high-frequency trading typically involves rapid cancellations and modifications of outstanding orders, an increase in high-frequency trading activity may be captured by an increase in the OTR.[18]

We proxy for order flow fragmentation by the dispersion of trading volume across trading venues. In particular, we define our measure of order flow fragmentation for each stock i on date t as the number of shares traded on venue v relative to the total trading volume across the OSE, CHI, TQ, and BATS. This measure can be interpreted as the daily market share of venue v in stock i.

3.4 Descriptives I: Stock liquidity at the OSE (2007–2009)

To place the tick size war of 2009 in a broader context, Figure 3 plots time-series of stock liquidity and stock prices for OBX index stocks at the Oslo Stock Exchange in the period 2007 to May, 2009. The figure shows that stock liquidity worsened significantly as stock prices declined during the financial crisis in the Autumn of 2008. During the first few months of 2009, however, both stock prices and stock liquidity at the OSE were gradually improving. This is particularly visible for average quoted spreads, which declined from 0.5% at the height of the financial crisis to about 0.25% in May, 2009 — almost the same level as before the crisis.

The sample period we consider surrounding the tick size war — the calendar year 2009 — is therefore in the tail-end of the financial crisis in 2008. This means that our data are drawn from a

[18] The OTR is also commonly referred to as the 'quote-to-trade' or the 'message-to-trade' ratio. Jørgensen et al. (2016) provide more details on order-to-trade ratios at the OSE.

period when stock liquidity at the Oslo Stock Exchange was improving for reasons that are likely to be unrelated to the tick size war of 2009. If unaccounted for in the empirical identification procedure, these pre-existing trends will erroneously be attributed to the estimated impact of the tick size war. In our empirical analysis of the impact of the tick size war on stock liquidity (Section 5), we attempt to overcome the problem of confounding pre-existing trends by using a difference-in-differences approach, which allows us to control for marked-wide trends in stock liquidity.[19]

3.5 Descriptives II: Trading at the OSE, Chi-X, Turquoise, and BATS (2009)

Table 2 summarizes our main outcome variables for the period January–May 2009 (the period before the tick size war) separately for the OSE, Chi X, BATS, and Turquoise. The table shows that the four stock exchanges in our sample differ notably in terms of estimated market quality. Transaction costs are smallest at the OSE with an average effective spread of 0.13%, followed by Turquoise with an average effective spread of 0.23%. The most expensive trading venue is Chi-X, with an average effective spread of 0.56%. Similarly, for our other two measures of transaction costs, relative and realized spreads, transaction costs are considerably smaller at the OSE than at the competing stock exchanges.[20]

The OSE order books are also by far the deepest. The average order book depth at the OSE is 733 thousand NOK. While this average to some degree is inflated by the depth in Statoil (The median OSE depth is 442 thousand), all the other exchanges (Chi-X, BATS, and Turquoise) have depths below 200 thousand. The OSE is also (by far) the most actively traded venue. Consequently, the OSE holds a commanding market position for trading in stocks with an OSE primary listing. The average market share of OSE in the period January–May 2009 is 99%. The Chi-X market share is 1.3% in the shares they offer trading in while BATS and Turquoise hold market shares of less than half a percent.

4 Market shares during the tick size war

In this section, we explore the impact of the tick size reductions during the tick size war in June 2009 on the distribution of market shares across stock exchanges. Consistent with theoretical predictions by Buti et al. (2015), we find that that small-tick size markets capture market shares from markets that keep large tick sizes. This finding suggests that competitive stock exchanges

[19] The Internet Appendix provides further descriptive statistics concerning the evolution of stock liquidity at the Oslo Stock Exchange, including summary statistics of our market quality measures both before and after the tick size war (in 2008 and 2010).

[20] Notice, however, that a direct comparison of transaction costs across exchanges may be misleading. For example, as indicated by the number of observations, Chi-X is active in more stocks than the other competing stock exchanges, BATS and Turquoise. That BATS and Turquoise appear to have smaller transaction costs than Chi-X may be because their trading activity is limited to only the most liquid stocks. Another reason to caution against a direct comparison of transaction costs is that our spread measures of liquidity do not account for the maker-taker fees applied at the MTFs. As such, we are comparing the gross transaction costs between venues, which may differ substantially from the net transaction costs, depending on the aggressiveness on the trading strategy.

may have an incentive to undercut other exchanges' tick sizes, as such tick size competition can allow them to increase their market share.

4.1 Results: Distribution of market shares

We begin our empirical analysis by exploring the evolution of market shares during the tick size war. To quantify the changes in market shares, we define three time periods. We define a *pre-war* period from May 1 to May 31, a *break-out* period from June 1 to July 5, and a *retaliation* period from July 6 to August 31. Within each of these time periods, we compute market shares for each stock i on date t for trading venue v.

In Table 3, we present the average market share for each trading venue in each of the three time periods, as well as the change in average market shares between a given time period and the pre-war period. The change in average market shares is obtained in a univariate regression framework where we compare daily observations of market shares in one period (the break-out period or the retaliation period) to daily observations of market shares in the pre-war period. We conduct a similar regression analysis for the natural logarithm of daily trading volume, to understand whether market share changes arise from flows of trading volume from one exchange to another, or alternatively from trader entry or exit.

Table 3 shows a considerable shift in market shares from the OSE to Chi-X. Before the tick size war, OSE market shares averaged 97.6% while Chi-X, the biggest competitor, operated with an average market share of 2.19%. During the break-out period, OSE market shares declined by a highly statistically significant 2.86 percentage points. These market shares were captured almost exclusively by Chi-X, which saw its market share more than double in the same period. Table 3 also shows that the shift in market shares appears to be driven by a flow of trading volume from the OSE to Chi-X — trading volume at the OSE fell by 26% after the Chi-X tick size reduction while trading volume at Chi-X increased by 68%. Turquoise market shares for OSE listed stocks increased slightly, while we find no impact on the market shares of BATS. Most of the order flow fragmentation occurs during the break-out period in June, while market shares remain relatively stable following the OSE tick size period (the retaliation period).[21]

To assess whether it is plausible that the market share changes in Table 3 are causally linked to tick size reductions, Figure 4 provides evidence on the timing of the market share changes. The figure shows an immediate and sizeable transfer of market shares from the OSE to Chi-X on the day of the Chi-X tick size reduction. Market shares for Turquoise and BATS show no such patterns. Following the OSE decision to reduce tick sizes in July, the OSE reclaims some of its lost market shares from Chi-X. Overall, Figure 4 provides appealing evidence that the

[21]Though our findings in Table 3 mostly pass the bar of statistical significance, it is not clear how we should assess the economic significance of the tick size reductions during the tick size war. On the one hand, a market share transfer of approximately 3% only amounts to a 50 million USD loss in trading volume, given a total trading volume of 10.22 billion NOK at the OSE on May 29, 2009. On the other hand, the 3% market share change was sufficient to prompt the OSE to make considerable changes to its market structure. It may be the case that the OSE judged the 3% market share change as economically sufficient by itself to respond to the Chi-X tick size reduction. More realistically, however, the OSE responded because the Chi-X tick size reduction also had an impact on the overall quality of trading at the OSE. In Section 5, we explore the market quality dimension of the tick size war.

market share changes during the Summer of 2009 are causally related to the tick size reductions during the tick size war.

Why did Chi-X, but not the other tick size-reducing exchanges, capture market shares from the OSE? The answer to this question is most likely a combination of three factors. First, Chi-X probably benefited from a 'first mover' advantage. Traders may have been settled and content with trading on the Chi-X platform when Turquoise and BATS decided to reduce their tick sizes. Second, out of the four stock exchanges, Chi-X operated with the smallest tick sizes during the *break-out* phase, meaning that the Turquoise and BATS tick size reductions offered nothing extra compared to Chi-X. Third, trading at Chi-X was already established and well-functioning before the tick size war; its market share, trading volume, and order book depth was reasonably high compared to Turquoise and BATS (see Table 2), which may explain why traders migrated to Chi-X and not the two other MTFs.

5 Tick size competition and market quality

Section 4 shows that the tick size war between the OSE, Chi-X, Turquoise, and BATS during the Summer of 2009 led to considerable shifts in the distribution of market shares across stock exchanges. In particular, the OSE experienced a considerable loss of market share to Chi-X. This section uses a difference-in-differences design to explore the impact of the tick size war on various measures of market quality.

5.1 Empirical specification

We use a difference-in-differences specification to estimate the impact of the tick size war on market quality. In our setting, the difference-in-differences approach involves comparing changes in market quality for a group of 'treated' stocks that were directly affected by the tick size reductions during the tick size war to changes in market quality in an unaffected 'control group' of stocks. This comparison between 'treated' and 'control' stocks is possible in our setting because only a subset of all OSE stocks were listed for trading at competing exchanges and therefore affected by the Chi-X, Turquoise, or BATS tick size reductions. The remaining OSE stocks were only traded at the OSE and were not affected by the tick size reductions (the 'control' group)

The most useful feature of the difference-in-differences design is that it allows us to control for confounding market-wide trends. This is achieved by estimating the effect of the tick size reductions during the Summer of 2009 net of the time trend in the control group of unaffected stocks. Controlling for market-wide trends is crucial in our setting since, as illustrated in Section 3.4, the sample period we consider is at the tail-end of a long positive trend in stock liquidity. If unaccounted for — using for instance a simple before-and-after event study design — this pre-existing trend would be attributed to our estimate of the impact of the tick size war on stock liquidity.[22]

[22] In the Internet Appendix we estimate before-and-after event study designs that do not account for pre-existing

We are mostly interested in the impact of the tick size war on market quality at the OSE and Chi-X, since trading at BATS and Turquoise appears to be largely unaffected by the tick size war. For this reason, we define two separate treatment groups that we evaluate in the difference-in-differences specification. The first treatment group is OBX index stocks traded on the OSE. The second treatment group is OBX index stocks traded on Chi-X. Both groups were directly affected by the Chi-X tick size reduction for OSE listed stocks on June 1, 2009 (labelled t_1^*) and the OSE tick size reduction for OBX index stocks on July 6, 2009 (labelled t_2^*).

We compare separately the evolution of stock trading in our two treatment samples to a single control sample. Our initial control sample consists of non-OBX index OSE stocks that were not traded on Chi-X, Turquoise, or BATS throughout the calendar year 2009. Since these stocks were not traded on any of the three MTFs, they were not directly affected by the MTF tick size reductions during June 2009. Moreover, since these stocks did not belong to the OBX index, they were not directly affected by the OSE tick size reduction on July 6. In order to maximize the comparability between our highly liquid OBX index treatment group stocks and our control group stocks, we use as control sample the 25 most-traded non-OBX stocks, where we use overall trading volume during May 2009 to rank the stocks outside the OBX index.

The difference-in-differences design is implemented with the following regression model:

$$y_{it} = \alpha_i + \alpha_t + \tau Treatment_{it} + \omega_{it}, \tag{1}$$

where $Treatment_{it} = 1$ for stock i that belongs to the treatment group on date $t \geq t^*$ and zero otherwise; α_i are stock-level fixed effects; and α_t are date-level fixed effects. The inclusion of stock and date fixed effects in equation (1) controls for fixed differences in y_{it} between treatment and control sample stocks and ensures that the effect of $Treatment_{it}$ on y_{it} is measured net of the time trend in the control sample. Under the identifying assumption that treatment and control stocks follow the same trend in y_{it} in the absence of treatment, the coefficient τ in equation (1) can be interpreted as the causal impact of the tick size war on stock market quality.

Equation (1) is estimated separately for the two tick size reduction events of interest — the Chi-X tick size reduction on June 1, 2009 (t_1^*) and the OSE tick size reduction on July 6, 2009 (t_2^*). We restrict the sample period surrounding the June 1 event to April 1 to July 5. Surrounding the July 6 event, we use a sample period from June 1 to August 31. Figure 5 illustrates how our sample periods are defined.

5.2 Results: Market quality

In the top panel of Table 4, we use the difference-in-differences specification to assess the impact of the Chi-X tick size reduction (t_1^* = June 1, 2009) on the quality of trading at the OSE and

trends. The Internet Appendix also provides further descriptive evidence for why such before-and-after designs are unlikely to inform us about the causal impact of the tick size war on stock liquidity. As an alternative way to estimate the impact of the tick size war on market quality, the Internet Appendix also includes an estimation of a so-called regression discontinuity design. The results from this specification are broadly consistent with the results we obtain with the difference-in-differences design.

Chi-X. The table shows that stock liquidity at the OSE deteriorates as a result of the Chi-X tick size reduction. For example, effective (realized) spreads increase by 9.9 (6.5) basis points for OSE listed stocks directly affected by the tick size reduction relative to a control group of OSE listed stocks not affected by the tick size reduction. These findings are robust to alternative specifications of the difference-in-differences design.[23] Despite capturing market shares, we find only weak evidence that Chi-X market quality increased. In particular, effective and realized spreads at Chi-X decrease but the effects are statistically insignificant. Order book depth at Chi-X improves by almost 15%, but the coefficient is only statistically significant at the 10% level.

To what extent are the observed changes in market quality at the OSE and Chi-X accounted for by a simple redistribution of trading volume? Section 4.1 documents that the shift in market shares between the OSE and Chi-X is mostly accounted for by a flow of trading volume from the OSE to Chi-X. Given the tight relationship between trading volume and measures of stock liquidity, one could imagine that the reduction (improvement) in stock liquidity at the OSE (Chi-X) is mechanically related to the flow of trading volume documented in Section 4.1. To determine the extent to which the observed changes to market quality are driven by changes in the distribution of trading volume, we include trading volume as a control variable in our difference-in-differences regressions. For Chi-X, we find that the entire increase in order book depth can be accounted for by an increase in trading volume. Meanwhile, for the OSE, the negative effects of the tick size war on trading costs persist even after controlling for trading volume.

In the bottom panel of Table 4, we evaluate the impact of the OSE tick size reduction (t_2^*=July 6, 2009) on stock market quality. The OSE tick size reduction causes a considerable reduction in order book depth at both the OSE (−42.5%) and Chi-X (−20%) — both effects measured relative to OSE listed stocks with no tick size change. We find no impact of the OSE tick size reduction on spread measures of liquidity at the OSE. In contrast, effective spreads at Chi-X appear to decline slightly following the OSE tick size reduction. Meanwhile, this effect fails to replicate in alternative specifications of the difference-in-differences design, which means that we cannot place much weight on this finding (see the Internet Appendix).[24]

6 Mechanisms

Section 4 shows that Chi-X was able to capture market shares from the OSE by reducing its tick size, while Section 5 shows that this tick size competition was detrimental to stock liquidity. In this section we investigate competing mechanisms that can potentially explain the redistribution

[23]The Internet Appendix estimates alternative specifications of the difference-in-differences design. For example, the Internet Appendix shows that our results are robust to alternative control groups and shorter sample periods.

[24]A future version of this paper will expand the empirical analysis in two ways. First, in addition to estimating the effects of the individual events of the tick size war, we will also estimate the overall impact of the tick size war. Second, we will include measures of price impact and price informativeness as outcome variables in our difference-in-differences estimations.

of market shares from the OSE to Chi-X following the Chi-X tick size reduction, and the ensuing changes to market quality.

We begin by distinguishing between two competing mechanisms for how between-exchange tick size differences affect the distribution of market shares. The first mechanism we consider is whether tick sizes affect market share changes by constraining the bid-ask spread in the main market, inducing traders to 'queue-jump' by sending orders to alternative markets where the bid-ask spread is less contrained (Buti et al., 2015). The second mechanism we consider is whether the tick size affects the distribution of market shares through its impact on the trading behavior of high-frequency traders (HFT).

6.1 Empirical specification

To distinguish between our two candidate mechanisms we estimate a cross-sectional regression where the change in the OSE market share is explained by proxies related to our candidate mechanisms. To implement this empirical test, we collapse all our stock characteristics into averages within two separate time periods — a pre-tick size war period in May 2009 and a post-tick size war period between June 1, 2009 and July 6, 2009 — and estimate the following cross-sectional regression:

$$\Delta^{Post-Pre} Marketshare_i^{OSE} = \alpha_0 + \beta \boldsymbol{X}_i^{Pre} + \varepsilon_i \tag{2}$$

where $\Delta^{Post-Pre} Marketshare_i^{OSE}$ is the change in average OSE market share between May 2009 (pre period) and June 2009 (post period) for stock i, and $\boldsymbol{X}_i$ is a vector of average pre-tick size war covariates, which includes proxies for tick size constraints and HFT at the OSE. The vector $\boldsymbol{X}_i$ is constructed using data from the pre-tick size war period to avoid that the stock characteristics in $\boldsymbol{X}_i$ themselves are affected by the tick size war. The sample we use to estimate equation (2) includes only the stocks that were directly affected by the Chi-X tick size reduction (our main sample of OBX index stocks).

Similar to our approach in Section 5, we simplify the analysis by focusing on the the trading that occurs at the OSE and Chi-X. Consequently, $Marketshare_i^{OSE}$ is computed as the distribution of share trading volume between the OSE and Chi-X.

6.2 Our proxies: 'Tick constrained' and 'Order-to-trade'

Before we estimate equation (2), we define our empirical proxies for the stock-level extents of tick size spread constraints and HFT activity.

First, an order book is potentially constrained by the tick size when the distance between the best bid and best ask is equal to a single tick. Tick size constraints can either be measured by a binary variable for whether or not the bid-ask spread is constrained by the tick size, or by a discrete variable that counts the number of ticks between the best bid and ask. We use a binary variable constructed following the procedure in O'Hara et al. (2015), where we first compute the average number of ticks-per-quoted-spread (the quoted spread divided by the tick size) during the pre-war period, and then define a stock to be *Tick Constrained* if the average number of ticks-per-quoted spread is less than two.[25]

Second, to proxy for the extent of HFT activity, we use the order-to-trade ratio (OTR). The OTR is a commonly-used proxy for HFT activity, and is computed by counting the number of orders that are submitted to the limit order book and dividing this count by the number of executed trades. Descriptive statistics of the OTR are presented in Table 2. In the period January–May 2009, the average number of orders per executed trade at the OSE was 8, with a standard deviation of 5.6 and a median of 6.4.

In addition to our proxies for tick size constraints and HFT activity, we allow for alternative determinants of market share changes. In the regressions we include measures of trading costs (quoted and effective spreads) as well as measures of trading volume (NOK volume or order book depth).

6.3 Results I: Cross-sectional regressions

Table 5 presents estimates of the cross-sectional regression (2), where the change in OSE market share between the pre-war period and the break-out period is the dependent variable. Starting with our proxy for tick size constraints, the table shows a negative regression coefficient, which indicates that tick size constrained shares fragment more as a result of the Chi-X tick size reduction in June 2009. This positive relationship between tick size constraints and order flow fragmentation is in line with theoretical predictions (Buti et al., 2015). However, the relationship is not statistically significant in any of the regression specifications in Table 5.

In contrast, we find a strong and statistically significant relationship between our measure of HFT in the pre-war period and subsequent order flow fragmentation. Specifically, stocks that have more HFT activity at the OSE fragment more following the Chi-X tick size reduction. This result remains statistically significant across various regression specification. As a consequence, our results more strongly favor that between-exchange tick size differences affect the distribution of market shares through its impact on HFT activity. This result contrasts with the existing empirical evidence from U.S. markets, which suggests that between-exchange tick size differences affect market shares because of queue jumping driven by differences in the severity of spread constraints (e.g. Buti et al. 2015).

One possible explanation for why we find no significant relationship between our measure of tick size spread constraints at the OSE and the change in market share is that our proxy may not capture the aspects of tick size spread constraints that are relevant for traders. After all, what matters to traders is not necessarily whether the bid-ask spread at the OSE is constrained by the minimum tick or not. Instead, traders may care about whether the spread at the OSE is more or less constrained than at Chi-X.

Another possible explanation for the lack of correlation is that our measure of tick size spread constraints is computed in the pre-tick size war period, and not during the tick size war. One

[25] In the period January–May 2009, the average ticks-per-spread has a minimum of 1.272 and a median of 1.866. In unreported regressions, we have also used the actual number of ‘ticks per spread’ as an explanatory variable, instead of the binary *Tick Constrained* variable. The conclusions with this specification are similar.

can imagine that Chi-X's tick size reduction relaxed the tick size spread constraints at Chi-X compared to the OSE, an aspect which we do not capture with our regression specification (2).

To address these potential concerns, we construct an alternative proxy which attempts to capture the difference in tick size spread constraints between the OSE and Chi-X. We also measure this difference during, instead of before, the tick size war. To this end we estimate the following cross-sectional regression model:

$$\Delta^{Post-Pre} Marketshare_i^{OSE} = \alpha_1 + \beta_1 \left(TS_i^{OSE} - TS_i^{CHI} \right)^{Post} + \beta_2 \left(\boldsymbol{X}_i^{OSE} - \boldsymbol{X}_i^{CHI} \right)^{Pre} + \varepsilon_i \tag{3}$$

where TS_i^{OSE} and TS_i^{CHI} are the average ticks-per-quoted-spread measured during the tick size war period at the OSE and Chi-X, respectively. To address the possibility that between exchange differences in tick size spread constraints are correlated with between-exchange differences in stock liquidity, we generate a set of *relative liquidity* measures. In particular, the vector term $\left(\boldsymbol{X}_i^{OSE} - \boldsymbol{X}_i^{CHI} \right)$ captures the differences in our measures of stock liquidity and trading volume at the OSE and Chi-X.

Table 6 presents estimates from our cross-sectional regressions using differences in trading characteristics between the OSE and Chi-X as explanatory variables. The table confirms our previous findings that the extent of tick size constraints is a poor explanatory variable for the extent of market share changes during the tick size war. The only statistically significant explanatory variable we find is trading volume — shares that tend to be heavily traded at the OSE compared to Chi-X fragment less despite cross-market differences in the tick size.

6.4 Results II: High-frequency trading

Section 6.3 shows that the stock-level change in OSE market share following the Chi-X tick size reduction is positively related to the stock-level extent of high-frequency trading (HFT) at the OSE, even after controlling for observable characteristics of the stock's trading environment. This result is consistent with HFTs routing their orders to small-tick exchanges rather than large-tick exchanges. In this section we further explore the potential mechanism that HFTs can account for the observed redistribution of market shares from the OSE to Chi-X.

Two data limitations force us to rely on indirect empirical evidence in support of the HFT mechanism. First, the ideal empirical test for whether HFTs account for the market share changes would be to explore whether HFT activity at Chi-X increased after its tick size reduction, and that HFT activity at the OSE decreased. Unfortunately, our data do not permit such a test. This is because we can only proxy for HFT activity for trading at the OSE, and not for trading at Chi-X. The second limitation is that aggregate HFT activity at the OSE is unlikely to change much on account of the three percentage point market share loss to Chi-X. Nevertheless, we proceed by shedding light on two mechanisms that can illustrate why HFTs may be important in our setting.

The first mechanism we consider is that between-exchange tick size differences can create mechanical price differences for the same security at different stock exchanges, which allows for profitable cross-market arbitraging. For instance, a coarser price grid at the OSE implies that it is more difficult to bid the security price to its marginal fundamental value than at Chi-X. Traders with access to both markets can exploit these arbitrary between-exchange price differences by, for example, buying shares at Chi-X for prices that are unattainable at the OSE, and selling the shares in the price-constrained OSE market. HFTs' speed advantage make them prime candidates for exploiting such cross-market arbitrage opportunities, because it is typically the first mover who gets the best prices.

We leverage the granularity of our data to assess whether the tick size war introduced cross-market arbitraging opportunities, and whether investors actually traded on these arbitrage opportunities. To do so, we divide the trading day into separate five-minute intervals (e.g., 09:00am – 09:05am, or 09:06am – 09:10am), and collect from each five-minute interval the highest and the lowest trade prices that occur at the OSE.[26] Next, we infer whether trades at Chi-X in the same five-minute intervals occur at prices that are within the price bands at the OSE. The idea behind this comparison is that the coarse price grid at the OSE may prevent trades from happening at certain prices, while the granular price grid at Chi-X can accommodate these trades. For this reason, more Chi-X trades happening outside the OSE price bands indicates that traders route their orders to Chi-X to achieve better prices, perhaps with the intention of offloading the position for profit at the OSE.

Figure 6 presents striking evidence that the tick size war between the OSE and Chi-X introduced between-exchange price differences that traders acted on. Before the June 1 Chi-X tick size reduction, nearly 90% of all trades at Chi-X took place within the price bands established at the OSE. However, during the tick size war a much larger portion of Chi-X trades took place at prices outside the OSE price bands. Indeed, immediately after the Chi-X tick size reduction, the fraction of Chi-X trades occurring outside the OSE price bands increased from approximately 10% to more than 20%. In other words, as prices at the OSE and Chi-X deviate more often, cross-market arbitraging becomes more profitable.

However, cross-market arbitraging cannot be the only mechanism to explain why Chi-X captures market shares from the OSE. For one, cross-market arbitraging is a two-sided trading strategy, which means that any position that HFTs take at Chi-X is offset by an equal-sized position at the OSE. This implies that trading volume should increase at both the OSE and Chi-X while market shares remain the same, which is at odds with our findings in Section 4.1. An alternative mechanism, which is also compatible with the findings in Figure 6, is that the small tick size at Chi-X reduces the marginal cost of undercutting existing quotes, which makes it easier and perhaps more profitable for HFTs to pick off stale quotes at the arrival of new information.

To assess whether it is plausible that HFTs are drawn to Chi-X because its small tick size makes it easier for HFTs to implement their trading strategies, we explore how HFTs at the

[26] By using five-minute intervals instead of, for instance, one-second or split second intervals, we at least partly circumvent the problem that trading feeds from different exchanges may be imperfectly time-synchronized. Differences in processing times across exchanges makes it difficult to determine whether trades that are reported within very small time periods actually took place at comparable times.

OSE responded to the OSE tick size reduction on July 6, 2009. Figure 2 plots our measure of HFT activity throughout the calendar year 2009, separately for OBX index stocks directly affected by the tick size war and for non-OBX index stocks unaffected by the tick size war. The most striking feature of Figure 2 is that HFT activity at the OSE increased notably for OBX index stocks after the OSE tick size reduction but remained stable for the unaffected non-OBX stocks. This finding suggests that HFTs trade more actively when tick sizes are small, which can help explain our finding in Section 6.3 that the extent of market share losses at the OSE during the tick size war positively correlates with HFT activity.

6.5 Discussion of potential mechanisms

Section 4 shows that Chi-X was able to capture market shares from the OSE by reducing its tick size. Moreover, Section 5 shows that trading costs at the OSE increased as a consequence of the Chi-X tick size reduction, while trading costs at Chi-X remained unchanged. The current section presents supplementary evidence that high-frequency traders (HFTs) appear to be responsible for the observed redistribution of market shares from the OSE to Chi-X. We now propose two potential mechanisms that can unify our findings on HFT trading strategies and trading costs at Chi-X and the OSE during the tick size war.

The first mechanism we have in mind is that HFTs route their orders to the markets that offer the smallest tick size, in our case Chi-X, which drives the observed redistribution of market shares from the OSE to Chi-X. Moreover, since we find that stock liquidity at Chi-X does not appear to improve from the inflow of HFT trading volume, we conclude that these HFTs consume liquidity and do not supply liquidity. Alternatively, the finding that HFT activity does not improve liquidity at Chi-X is consistent with HFTs being informed investors whose trading imposes an adverse selection cost for limit order traders at Chi-X, which forces bid-ask spreads to widen.

Meanwhile, the interpretation that informed and liquidity-demanding HFTs migrate the OSE in favor of Chi-X cannot explain why trading trading costs at the OSE increased after the Chi-X tick size reduction. This is because trading costs at the OSE should worsen when informed liquidity-demanders leave the exchange. We interpret the finding that the same trading volume can have opposite impacts on the trading costs at the OSE and Chi-X as consistent with a mechanism where HFTs switch from trading as liquidity-providers in the large-tick OSE market to trading as liquidity-demanders in the small-tick Chi-X market.

Most of our empirical analysis has focused on HFT trading strategies and investors' trading costs following the June 1 Chi-X tick size reduction, and relatively little attention has been given to the OSE retaliatory tick size reduction in July, 2009. This is mainly because most of the change in market shares during the tick size war occurs in a small time period following the Chi-X tick size reduction (see Section 4.1). However, though we find little change in market shares following the OSE retaliatory tick size reduction, we do observe that investors adapt their trading strategies to the new tick size. For example, we find that order book depths at both the OSE and Chi-X declined considerably as the OSE tick sizes came down. This finding suggests that tick size reductions lower the incentives to post limit orders at the top of the order book. Moreover, this finding suggests that tick size reductions in one market can have negative spill-over effects on the order book depths in markets that do not change their tick size.

7 Conclusion

This paper studies a situation where competition can induce stock exchanges to implement market design changes that worsen trading conditions for market participants. Our empirical analysis considers an event in 2009 where three European stock exchanges, Chi-X, Turquoise, BATS Europe, reduced their tick sizes (the smallest price increment on the exchange) for stocks with an Oslo Stock Exchange (OSE) primary listing. The OSE quickly responded by reducing its own tick sizes, before all the exchanges agreed on a common tick size structure. We find that the tick size-reducing exchanges captured market share from the OSE, and that the competitive tick size reductions increased trading costs for market participants. High frequency trading appears to be the main driver behind the market share and trading cost results.

The results in this paper contribute to the existing empirical literature on tick sizes. First, a recent literature shows that trading venues that offer small tick sizes can capture market shares from large-tick trading venues (e.g. Bartlett and McCrary 2015, Biais, Bisière, and Spatt 2010, Kwan et al. 2015). Consistent with the existing literature, we find that trading platforms with relatively small tick sizes capture market share from large-tick trading platforms. We add to the existing literature by exploring the tick size that arises endogenously through competition between stock exchanges that strategically adjust their tick size, and estimate the effects of this competitive tick size on market quality.

Second, our results connect to the empirical debate over HFTs' optimal response to tick size changes. O'Hara et al. (2015) and Yao and Ye (2015) argue that HFTs become more active in liquidity provision and have larger profit margins in a large-tick environment. They propose that HFTs' speed advantage becomes more valuable when price competition is constrained by the tick size. Our results, in contrast, suggest that HFT seem to migrate large-tick exchanges in favor of small-tick exchanges. The conflicting results in our paper can indicate that certain types of HFT strategies may require a fine pricing grid whereas other HFT strategies, such as liquidity-provision, can benefit from a large tick size.

Finally, this paper provides empirical support for current market regulations in the United States that enforce a common tick size across competing stock exchanges, and for proposed regulations in Europe that aim to accomplish the same. Our results suggest that individual stock exchanges have an incentive to reduce the tick size to capture market shares and, at the same time, that competitive tick size reductions can reduce overall market quality. Policy makers can limit stock exchanges' ability to engage in such destructive tick size competition by strictly enforcing a shared tick size regime across all trading venues competing for the same order flow.

References

Bartlett, III, R. P. and J. McCrary (2015): "Dark Trading at the Midpoint: Pricing Rules, Order Flow and High Frequency Liquidity Provision," Working Paper, University of Berkeley.

BATS (2009): "Pan European Tick Size Pilot," BATS Trading Working Paper.

BATS-CHIX (2012): "BATS Chi-X Europe Reference Data Specification," Manual.

Biais, B., C. Bisière, and C. Spatt (2010): "Imperfect Competition in Financial Markets: An Empirical Study of Island and Nasdaq," *Management Science*, 56, 2237–2250.

Buti, S., F. Consonni, B. Rindi, Y. Wen, and I. Werner (2015): "Sub-Penny and Queue-Jumping," Working paper, Ohio State University, Columbus OH.

Cordella, T. and T. Foucault (1999): "Minimum Price Variations, Time Priority, and Quote Dynamics," *Journal of Financial Intermediation*, 8, 141–173.

Federation of European Securities Exchanges (2009): "European exchanges introduce harmonised tick size regimes in Europe," Press Release, Brussels, 30th June 2009.

Foucault, T., O. Kadan, and E. Kandel (2005): "Limit Order Book as a Market for Liquidity," *The Review of Financial Studies*, 18, 1171.

Glosten, L. R. and P. R. Milgrom (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of Financial Economics*, 14, 71 – 100.

Grant, J. (2009): "LSE bows to tick size pressure as war erupts," Financial Times, London. Available at https://www.ft.com/content/1b386f66-5b69-11de-be3f-00144feabdc0.

——— (2011): "Euronext sparks outrage with tick size reduction," Financial Times, London.

Hagströmer, B. and L. Nordén (2013): "The Diversity of High Frequency Traders," *Journal of Financial Markets*, 16, 741–770.

Holden, C. W., S. Jacobsen, and A. Subrahmanyam (2013): "The Empirical Analysis of Liquidity," *Foundations and Trends in Finance*, 8, 263–365.

Jørgensen, K., J. Skjeltorp, and B. A. Ødegaard (2016): "Throttling hyperactive robots - Order to Trade Ratios at the Oslo Stock Exchange," Working Paper, University of Stavanger.

Kwan, A., R. Masulis, and T. McInish (2015): "Trading Rules, Competition for Order Flow and Market Fragmentation," *Journal of Financial Economics*, 28, 592–636.

Kyle, A. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–35.

Meling, T. G. (2016): "Anonymous trading in equities," Working Paper, SSRN.

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

——— (2016): "The Economics of High-Frequency Trading: Taking Stock," *Annual Review of Financial Economics*, 8, 1–24.

O'Hara, M., G. Saar, and Z. Zho (2015): "Relative Tick Size and the Trading Environment," Working Paper, SSRN.

Pagano, M. (1989): "Trading Volume and Asset Liquidity," *Quarterly Journal of Economics*, 104, 255–274.

Securities and Exchange Commission (2012): "Report to Congress on Decimalization," SEC Report.

Verousis, T., P. Perotti, and G. Sermpinis (2017): "One size fits all? High frequency trading, tick size changes and the implications for exchanges: Market quality and market structure considerations," *Review of Quantitative Finance and Accounting*, forthcoming.

Yao, C. and M. Ye (2015): "Why Trading Speed Matters: A tale of Queue Rationing under Price Controls," Working Paper.

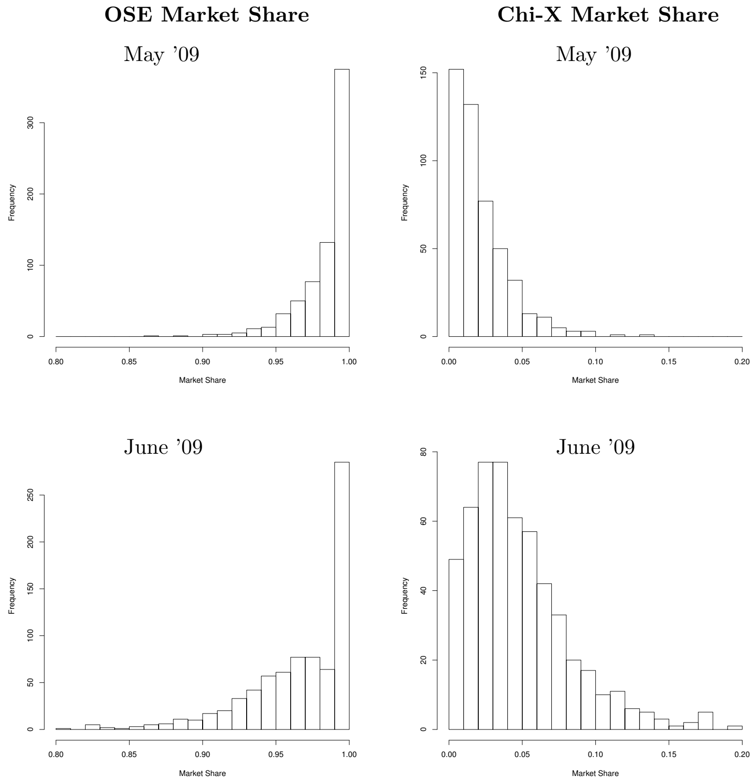

The figure presents the distribution of daily market shares at the Oslo Stock Exchange (left) and Chi-X (right). The top panel presents the distribution of market shares during May, 2009. The bottom panel presents the distribution of market shares during June, 2009.

Figure 1: Distribution of market shares, May-June 2009

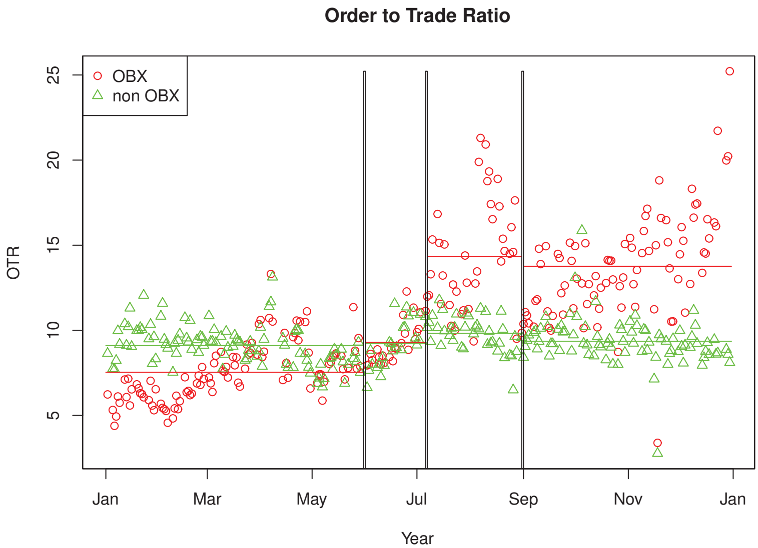

The figure presents daily cross-sectional averages of the order-to-trade ratio throughout the calendar year 2009, separately for OBX index stocks (red) and non-OBX index stocks (green). The left vertical break indicates June 1, 2009, the date when Chi-X reduced its tick size for OSE listed stocks. The middle vertical break indicates July 6, 2009, the date when the OSE reduced its tick size for OSE listed stocks. The right vertical break indicates August 31, 2009, the date when the OSE, Chi-X, Turquoise, and BATS Europe agreed on a common tick size for OSE listed stocks. Horizontal red and green line represent the average order-to-trade ratio within each sample window, for OBX and non-OBX index stocks, respectively.

Figure 2: Order-to-trade ratios at the OSE

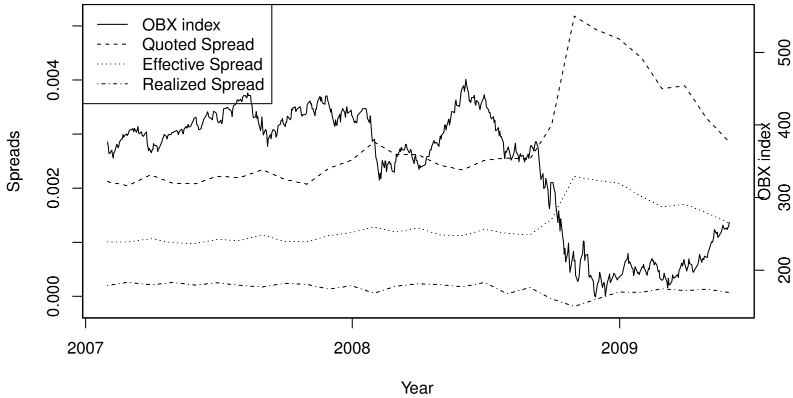

The figure presents the daily average price level of stocks in the OBX index (right axis) and monthly averages of three different spread measures of stock liquidity (left axis). The spread measures of liquidity are relative quoted spreads, effective spreads, and realized spreads (defined in Section 3.3). Our spread measures of liquidity are first computed on the stock-day level before they are averaged across all stocks in the OBX index on a monthly basis.

Figure 3: Stock prices and stock liquidity, 2007–2009

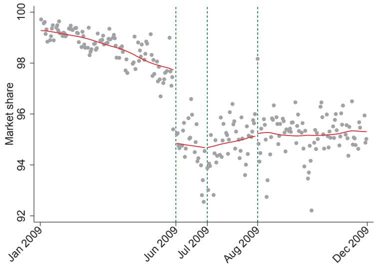

(a) OSE

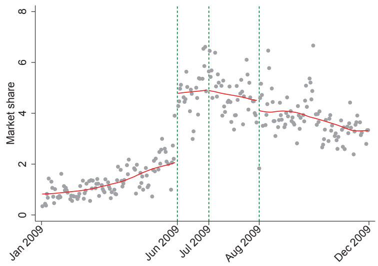

(b) CHI

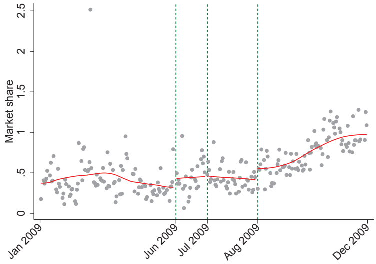

(c) TQ

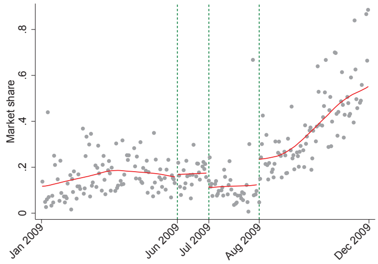

(d) BATS

The figure presents daily averages of stock-level market shares of trading in stocks with an Oslo Stock Exchange primary listing, presented separately for the Oslo Stock Exchange (OSE), Chi-X (CHI), Turquoise (TQ), and BATS Europe (BS). The market share in stock i on date t for venue v is given by the share trading volume on venue v relative to the share trading volume across OSE, Chi-X, Turquoise, and BATS. The left vertical break indicates June 1, 2009, the date when Chi-X reduced its tick size for OSE listed stocks. The right vertical break indicates July 6, 2009, the date when the OSE reduced its tick size for OSE listed stocks. Red lines are local polynomial smoothing regressions with a bandwidth of twenty trading days, that are fit separately within each of the sample windows.

Figure 4: Market shares throughout 2009

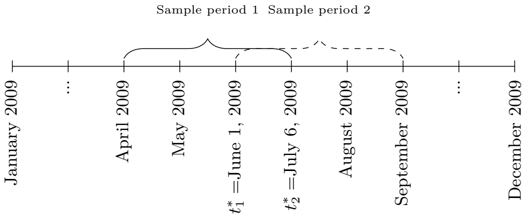

The figure illustrates how we define the difference-in-differences sample periods surrounding our two event dates. Our first event, t_1^*, is the beginning of the 'tick size war' on June 1, 2009. Our second event, t_2^*, is the Oslo Stock Exchange tick size reduction on July 6, 2009. First, surrounding the June 1, 2009, event, we restrict the sample period to April 1, 2009, to July 5, 2009. Second, surrounding the July 6, 2009 event, we restrict the sample period to June 1, 2009, to August 31, 2009. Solid curly braces span the sample period surrounding June 1, 2009. Dashed curly braces span the sample period surrounding July 6, 2009.

Figure 5: Illustration: Sample restrictions

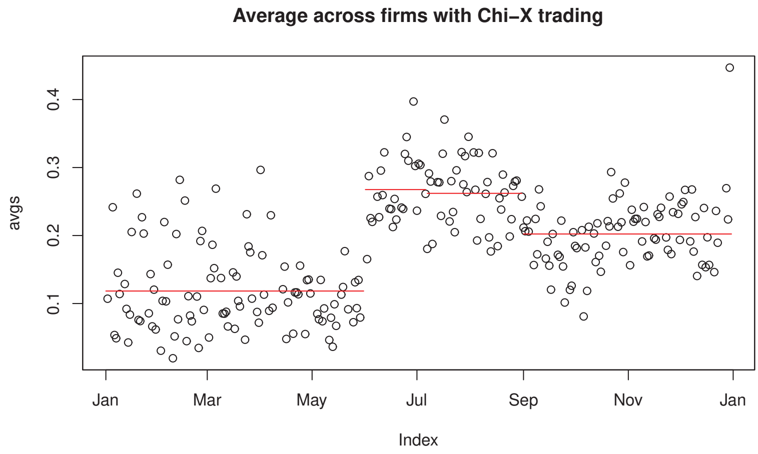

The figure presents daily cross-sectional averages of the fraction of Chi-X trades that take place outside the price bands established at the OSE. We define price bands by splitting the trading day into separate five-minute trading intervals (e.g., 09:00am – 09:05am, or 09:06am – 09:10am), and collect from each five-minute interval the highest and the lowest trade prices that occur at the OSE. Next, we infer whether trades at Chi-X in the same five-minute intervals occur at prices that are within the price bands at the OSE. We generate first a stock-day level variable which captures the fraction of trades at Chi-X that take place outside the OSE price bands, before we average this variable across all OSE stocks with trading at Chi-X.

Figure 6: Fractions of Chi-X trades outside the OSE price bands

Panel A: The Oslo Stock Exchange

– July 2009

	Price band			Tick Size
Most Liquid stocks (Statoil)		–		0.01
Other OBX stocks		–	14.99	0.01
	15	–	49.95	0.05
	50	–	99.90	0.10
	100	–	249.75	0.25
	250	–	499.50	0.50
	500	–		1.00
Non-OBX stocks (illiquid)		–	9.99	0.01
	10	–	14.95	0.05
	15	–	49.90	0.10
	50	–	99.75	0.25
	100	–	249.50	0.50
	250	–	-	1.00

July 2009

	Price band	Tick Size
All OBX Stocks	–	0.01

Fall 2009 –

	Price band			Tick Size
All OBX stocks		–	0.4999	0.0001
	0.5	–	0.9995	0.0005
	1	–	4.9990	0.001
	5	–	9.995	0.005
	10	–	49.990	0.01
	50	–	99.95	0.05
	100	–	499.90	0.1
	500	–	999.50	0.5
	1,000	–	4,999.00	1
	5,000	–	9,995.00	5
	10,000	–	–	10

Panel B: Chi-X and Turquoise/BATS

Chi-X – June 2009

	Price band			Tick Size
OBX Shares (selected)	0	–	9.99	0.001
	10	–		0.005

Turqoise/BATS – June 2009

	Price band			Tick Size
OBX shares (selected)		–	0.9999	0.0001
	1	–	4.9995	0.0005
	5	–	9.999	0.001
	10	–	49.995	0.005
	50	–	99.99	0.01
	100	–	499.95	0.05
	500	–	999.90	0.1
	1,000	–	4,999.50	0.5
	5,000	–	9,999	1
	10,000	–	99,995	5
	100,000	-		10

The table presents the tick size schedules used by the Oslo Stock Exchange (OSE), Chi-X, Turquoise, and BATS Europe during the tick size war of June, 2009. Chi-X implemented its tick size schedule on June 1, 2009, Turquoise on June 8, 2009, and finally BATS Europe on June 15, 2009. The tick size schedules for BATS Europe and Turquoise have been collected from BATS (2009). The tick size schedule for Chi-X has been collected from BATS-CHIX (2012) (the 'eurozone' tick size schedule).

Table 1: Tick size schedules at the OSE, Chi-X, BATS, and TQ.

	mean	std	min	median	max	n
Oslo Stock Exchange						
Relative spread (%)	0.404	0.212	0.089	0.341	1.668	2626
Effective spread (%)	0.132	0.065	0.036	0.116	0.573	2626
Realized spread (%)	0.025	0.061	-0.596	0.021	0.762	2626
Price Impact (%)	0.103	0.080	-0.234	0.090	0.923	2626
Depth (thousand NOK)	733	835	72	442	16758	2626
Realized Volatility (%)	0.970	1.790	0.179	0.677	46.864	2626
Volume (thousands NOK)	193023	322364	3000	71233	3942873	2626
Order to Trade Ratio	8.0	5.6	2.2	6.4	111.2	2525
Chi-X						
Relative spread (%)	2.366	1.705	0.159	1.809	8.589	2368
Effective spread (%)	0.556	0.437	0.059	0.414	3.248	1863
Realized spread (%)	0.174	0.513	-4.160	0.077	5.404	1859
Price Impact (%)	0.378	0.469	-4.077	0.292	3.760	1857
Depth (thousand NOK)	187	106	12	174	981	2388
Realized Volatility (%)	0.558	0.269	0.047	0.515	5.603	1693
Volume (thousands NOK)	2364	4594	0	782	66823	2507
BATS						
Relative spread (%)	0.696	0.752	0.099	0.529	9.856	1429
Effective spread (%)	0.294	0.281	0.042	0.219	4.209	654
Realized spread (%)	0.106	0.676	-7.046	0.113	4.043	653
Price Impact (%)	0.235	0.610	-2.375	0.157	7.383	629
Depth (thousand NOK)	78	45	16	74	993	1674
Realized Volatility (%)	0.500	0.305	0.045	0.434	3.033	415
Volume (thousands NOK)	212	363	1	93	5777	1581
TRQ						
Relative spread (%)	0.536	0.723	0.118	0.360	7.798	656
Effective spread (%)	0.233	0.265	0.047	0.172	3.155	608
Realized spread (%)	0.105	0.311	-1.751	0.073	2.251	611
Price Impact (%)	0.157	0.311	-1.458	0.104	2.280	599
Depth (thousand NOK)	136	71	3	124	801	750
Realized Volatility (%)	0.522	0.258	0.086	0.472	2.390	611
Volume (thousands NOK)	1618	2519	1	843	37203	889
Market Shares						
OSE	99.0	1.5	77.4	99.6	100.0	3747
Chi-X	1.3	1.5	0.0	0.8	22.6	2321
BATS	0.2	0.3	0.0	0.1	2.8	1613
TRQ	0.4	0.6	0.0	0.3	11.0	908

The table summarizes stock trading characteristics separately for trading at the Oslo Stock Exchange, Chi-X, BATS, and Turquoise. The sample period is January–May, 2009 (time period before the tick size war). Market quality measures: *Quoted (relative) spread*: The difference between the best bid and best ask in the order book, divided by price. Averaged across all order books during a trading day. *Effective spread*: Difference between trade price and a pre-trade benchmark, relative to benchmark. *Realized spread*: Difference between trade price and a post-trade benchmark, relative to trade price. *Price Impact*: Difference between post-trade and pre-trade benchmark, relative to pre-trade benchmark. *Depth*: The total (NOK) amount outstanding at the best bid and ask. *Volume*: The total amount (in NOK) traded. *Realized volatility*: The (uncentered) standard deviation over ten minute interval returns. *Order to Trade Ratio*: Ratio of messages to the exchange's order book divided by the number of consummated trades, on a daily basis. Only calculated for the OSE. *Market shares*: The proportion of share trading volume on a given trading venue relative to the total share trading volume across the OSE, Chi-X, BATS, and Turquoise. At the OSE, the sample comprises all OBX index stocks.

Table 2: Descriptive statistics, January–May 2009

	Pre-war May 1 - May 31	Break-out June 1 - July 5	Retaliation July 6 - August 31
Oslo Stock Exchange			
Market share	97.60	94.74	94.95
Diff.		-2.86***	-2.65***
Trading volume (log)	18.72	18.46	18.18
Diff.		-0.26***	-0.53***
Chi-X			
Market share	2.19	4.87	4.67
Diff.		2.67***	2.48***
Trading volume (log)	14.50	15.18	14.90
Diff.		0.68***	0.40***
Turquoise			
Market share	0.34	0.45	0.44
Diff.		0.11**	0.10**
Trading volume (log)	12.71	12.59	12.37
Diff.		-0.11	-0.34*
BATS			
Market share	0.16	0.17	0.11
Diff.		0.01	-0.04**
Trading volume (log)	11.74	11.52	10.57
Diff.		-0.22**	-1.17***

The table presents average market shares and trading volume for trading in stocks with an Oslo Stock Exchange primary listing, separately for the Oslo Stock Exchange, Chi-X, Turquoise, and BATS Europe. Market share in stock i on date t for venue v, is given by the share trading volume on venue v relative to the share trading volume across OSE, Chi-X, Turquoise, and BATS. Average market shares and trading volume are computed for three time periods: the pre-war period (May 1 to May 31); the break-out period (June 1 to July 5); and the retaliation period (July 6 to August 31). The table also presents the change in market share and trading volume between a given period (the break-out period or the retaliation period) and the pre-war period. The between-period changes in market share and trading volume are obtained by separately comparing daily observations of market shares or trading volume in either the break-out period or the retaliation period to daily observations of market shares or trading volume in the pre-war period in a regression framework. Standard errors are clustered at the stock level.

Table 3: Distribution of market shares during tick size war

Panel A: Chi-X tick size reduction ($t^* =$ June 1, 2009)

	Effective spread		*Realized spread*		*Depth*		*Volatility*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OSE								
τ	0.099***	0.091**	0.065**	0.056**	0.101	0.119	-0.001	-0.001
	(2.70)	(2.61)	(2.56)	(2.31)	(1.23)	(1.53)	(-0.68)	(-0.49)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	3018	3018	3021	3021	3157	3125	2921	2921
Adj. R^2	0.66	0.67	0.29	0.30	0.85	0.87	0.05	0.07
CHI								
τ	-0.052	-0.001	-0.008	0.030	0.148*	0.038	-0.000	-0.002**
	(-1.09)	(-0.01)	(-0.18)	(0.69)	(1.76)	(0.44)	(-0.51)	(-2.04)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	2825	2825	2825	2825	3106	3021	2629	2629
Adj. R^2	0.52	0.54	0.19	0.20	0.63	0.66	0.10	0.13

Panel B: OSE tick size reduction ($t^* =$ July 6, 2009)

	Effective spread		*Realized spread*		*Depth*		*Volatility*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
OSE								
τ	-0.002	0.004	-0.035	-0.026	-0.425***	-0.436***	-0.000	-0.000
	(-0.09)	(0.13)	(-1.66)	(-1.33)	(-5.23)	(-5.33)	(-0.16)	(-0.39)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	3121	3121	3120	3120	3271	3226	3022	3022
Adj. R^2	0.72	0.73	0.38	0.39	0.82	0.84	0.05	0.07
CHI								
τ	-0.077**	-0.070**	-0.023	-0.015	-0.201**	-0.207**	0.001	0.000
	(-2.21)	(-2.07)	(-0.92)	(-0.61)	(-2.44)	(-2.54)	(1.11)	(0.84)
Controls	No	Yes	No	Yes	No	Yes	No	Yes
N	3078	3078	3077	3077	3258	3200	2923	2923
Adj. R^2	0.63	0.64	0.30	0.32	0.58	0.61	0.12	0.14

The table presents estimates of τ from the difference-in-differences specification applied separately to the Chi-X tick size reduction ($t^* =$ June 1, 2009) and the OSE tick size reduction ($t^* =$ July 6, 2009). Surrounding the June 1, 2009, event, we restrict the sample period to April 1, 2009, to July 5, 2009. Surrounding the July 6, 2009 event, we restrict the sample period to June 1, 2009, to August 31, 2009. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Treatment_{it} + \omega_{it}$. $Treatment_t$ is a dummy variable equal to 1 for all treatment group observations on dates $t \geq t^*$. The difference-in-differences specification is estimated separately for two treatment groups. The first treatment group is OBX index stocks traded on the OSE. The second treatment group is OBX index stocks traded on Chi-X. Our control sample of stocks consists of the 25 most-traded (based on total trading volume) non-OBX index OSE stocks that were not traded on the multilateral trading facilities (MTFs) Chi-X, Turquoise, or BATS throughout the calendar year 2009. *Espread* is the effective spreads, in percentage points. *Rspread* is the realized spreads, in percentage points. *Depth* is order book depth, transformed with the natural logarithm. *Volatility* is measured in percentage points. Standard errors are clustered at the stock-level.

Table 4: Difference-in-differences

	Dependent variable:							
	Change in OSE Market Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Quoted (rel) spread	2.756						6.004	
	(2.625)						(4.073)	
Effective spread		13.538^{**}						18.728^{**}
		(6.456)						(7.161)
Depth			-0.001					0.001
			(0.005)					(0.006)
Volume				0.001			-0.00000	
				(0.003)			(0.005)	
Tick Constrained					-0.007		-0.005	-0.005
					(0.006)		(0.005)	(0.005)
Order to Trade						-0.002^{**}	-0.003^{**}	-0.002^{***}
						(0.001)	(0.001)	(0.001)
Constant	-0.034^{***}	-0.045^{***}	-0.010	-0.046	-0.022^{***}	-0.013^{*}	-0.019	-0.045
	(0.008)	(0.009)	(0.066)	(0.053)	(0.004)	(0.007)	(0.116)	(0.090)
Observations	26	26	26	26	26	26	26	26
Adjusted R^2	0.004	0.120	-0.039	-0.035	0.031	0.144	0.316	0.416
Note:								*p<0.1; **p<0.05; ***p<0.01

The table presents coefficient estimates from the regression $\Delta^{Post-Pre} Marketshare_i^{OSE} = \alpha_0 + \beta \boldsymbol{X}_i^{Pre} + \varepsilon_i$. The regressions explain changes in OSE market share between the pre-war period period (May '09) and the break-out period (June-July 5). Each column is a separate regression. Explanatory variables: *Quoted (relative) spread*: The difference between best bid and best ask in the order book, divided by price. Averaged across all order books during a trading day. *Effective spread*: Difference between trade price and a pre-trade benchmark, relative to trade price. *Depth*: The natural log of the total (NOK) amount outstanding at the best bid and ask. Volume: The natural log of the total amount (in NOK) traded. *Tick Constrained*: Dummy variable equal to one if "Spreads per tick" (Quoted spread divided by tick size) is less than two. *Order to Trade Ratio*: The number of orders (messages) in the trading system per trade. All the explanatory variables are measured as averages over daily observations at the OSE during May, 2009 (the pre-war period).

Table 5: Explaining market share changes with OSE characteristics

	Dependent variable:						
	Change in OSE Market Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Quoted (rel) spread	-0.156					-0.273	
	(0.274)					(0.313)	
Effective spread		-0.587					0.793
		(1.172)					(1.073)
Depth			0.002			0.003	
			(0.007)			(0.008)	
Volume				0.009^{***}			0.011^{***}
				(0.003)			(0.003)
Relative Tick					-5.861	-5.512	-5.698
					(6.370)	(6.670)	(5.078)
Constant	-0.029^{***}	-0.029^{***}	-0.030^{***}	-0.068^{***}	-0.019^{**}	-0.028	-0.064^{***}
	(0.004)	(0.004)	(0.011)	(0.011)	(0.008)	(0.017)	(0.013)
Observations	25	25	25	25	26	25	25
Adjusted R^2	-0.029	-0.032	-0.039	0.341	-0.006	-0.074	0.354
Note:					*p<0.1; **p<0.05; ***p<0.01		

The table shows results of estimation of
$\Delta^{Post-Pre} Marketshare_i^{OSE} = \alpha_1 + \beta_1 \left(TS_i^{OSE} - TS_i^{CHI}\right)^{Post} + \beta_2 \left(\boldsymbol{X}_i^{OSE} - \boldsymbol{X}_i^{CHI}\right)^{Pre} + \varepsilon_i$ The regressions explain changes in OSE market share between the pre-war period period (May, 2009) and the break-out period (June-July 5). Each column is a separate regression. Explanatory variables: *Quoted (relative) spread*: The difference between best bid and best ask in the order book, divided by price. Averaged across all order books during a trading day. *Effective spread*: Difference between trade price and a pre-trade benchmark, relative to trade price. *Depth*: The natural log of the total (NOK) amount outstanding at the best bid and ask. Volume: The natural log of the total amount (in NOK) traded. *Relative tick*: Quoted spread (average through trading day) divided by tick size. All explanatory variables X except *Relative tick* are first measured on a daily basis as $X_{OSE} - X_{CHI}$, and then averaged within the pre-war period May, 2009. *Relative tick* is measured as the daily difference in 'Spreads per tick' at the OSE and Chi-X, and averaged within the break-out period June 1 to July 6.

Table 6: Explaining fragmentation with difference main market (OSE) and aggressor (Chi-X)

Internet Appendix to Tick Size Competition, High Frequency Trading, and Market Quality

Tom Grimstvedt Meling and Bernt Arne Ødegaard

May 2017

A Causal identification of liquidity effects

Section 5 uses a so-called difference-in-differences design in an attempt to capture the causal impact of the tick size war for OSE listed stocks on stock liquidity at the OSE and Chi-X. The difference-in-differences specification in Section 5 uses a control sample of stocks that were not affected by the tick size war to account for market-wide confounding trends in stock liquidity. The purpose of this section is to illustrate that a simple before-and-after design, which does not account for market-wide trends in stock liquidity, is unlikely to capture the causal effect of the tick size war on stock liquidity.

A.1 The trouble with before-and-after designs

The purpose of our empirical tests in Section 5 is to capture the causal effect of the tick size reductions during the tick size war on stock liquidity at the OSE and Chi-X. Causal effect estimates can only be obtained by comparing the evolution (or level) of outcomes to a valid counter-factual — that is, the evolution of the same outcome in the absence of some 'treatment.' In our setting, we are interested in the counter-factual scenario of how measures of stock liquidity at the OSE and Chi-X would have evolved during June and July 2009 without the tick size war. In Section 5, we propose that the evolution of stock liquidity in stocks that were not affected by the tick size war provide a valid counter-factual for the evolution of stock liquidity in stocks that were directly affected by the tick size war.

An alternative empirical strategy, which is much-used in the market microstructure literature, is to compare outcomes after an event (for example, a tick size reduction) to the same outcome before the event — in effect treating the pre-event period as the counter-factual scenario. In Table A.1, we perform such a before-and-after analysis, and present averages of various market quality measures for four different time periods: the pre-war period (May, 2009); the break-out period (June 1 to July 5); the retaliation period (July 6 to August 31); and the post-war period (September). The before-and-after exercise indicates that spread measures of stock liquidity improve at both the OSE and Chi-X during the tick size war, while order book depths remain largely unchanged throughout June before plummeting in July.

There are (at least) two reasons why, in our case, before-and-after estimates are unlikely to inform us about the causal impact of the events of the tick size war on market quality. First, beginning in early 2009, measures of stock liquidity at the OSE were steadily improving for reasons unrelated to the tick size war (see Figure A.1). Going further back in time, as we show in Figure 3 in Section 3.4, we find that the persistent trends to stock liquidity in 2009 reflect

Liq.measure	Market	Pre May	Breakout 1jun-5jul	Retaliation 6jul-31aug	Post Sep
Quoted (relative) Spread (%)	OSE	0.308	0.302	0.292	0.239
	Chi-X	1.542	0.988	0.681	0.501
Effective Spread (%)	OSE	0.106	0.101	0.083	0.073
	Chi-X	0.388	0.320	0.229	0.157
Realized Spread (%)	OSE	0.018	0.017	0.015	0.017
	Chi-X	0.051	0.038	0.043	0.029
Realized Volatility (%)	OSE	0.909	0.779	0.591	0.553
	Chi-X	0.553	0.515	0.400	0.401
Depth (Thousand NOK)	OSE	983	951	564	734
	Chi-X	217	248	183	185
Volume (Million NOK)	OSE	280	212	149	208
	Chi-X	5	8	6	7
Turnover (%)	OSE	1.44	1.04	0.67	0.87
	Chi-X	0.02	0.04	0.03	0.03

The table reports subperiod averages of a number of market quality measures at the OSE and Chi-X. Market quality measures: *Quoted (relative) spread*: The difference between best bid and best ask in the order book, divided by price. Averaged across all order books during a trading day. *Effective spread*: Difference between trade price and a pre-trade benchmark, relative to trade price. *Depth*: The total (NOK) amount outstanding at the best bid and ask. Volume: The total amount (in NOK) traded. All the variables are measured as averages (panel A) over daily observations in the given time interval. The numbers are averages over stock-level averages.

Table A.1: Before-and-after estimates

a recovery from a period of low liquidity during the financial crisis in the Autumn of 2008. If unaccounted for, the before-and-after estimators assign such pre-existing trends to the impact of the tick size war.

The second reason why comparing market quality during June and July to market quality in May is unlikely to identify the causal impact of the tick size war, is that trading behavior tends to be different during the Summer months (June and July) on account of public holidays. To provide some perspective on this potentially confounding factor, in Table A.2, we present market quality statistics from the same subperiods as in Table A.1 but, instead, one year before (2008) and one year after (2010) the tick size war. Most notable is the tendency of trading volume to be considerably lower during the Summer months compared to both May and September.

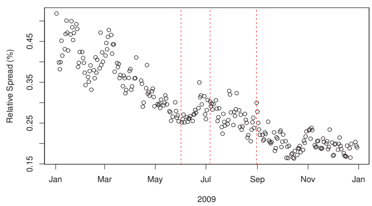

(a) Relative (Quoted) Spread

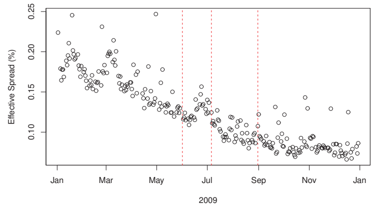

(b) Effective spread

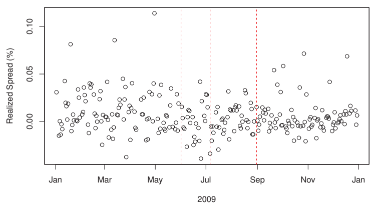

(c) Realized spread

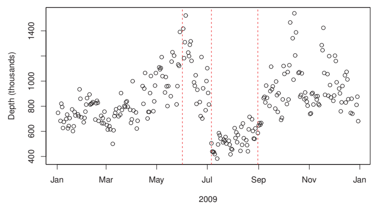

(d) Depth

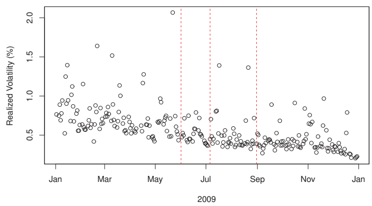

(e) Realized Volatility

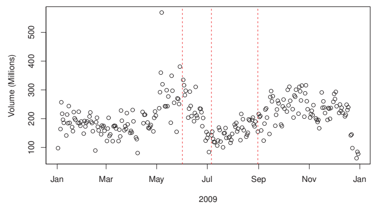

(f) Volume(NOK)

The figure presents time-series of market quality measures at the Oslo Stock Exchange, in the period January 1, 2009, to December 31, 2009. All observations are daily cross-sectional averages, computed across all OBX listed stocks. Panel (a) shows the effective spread, expressed in basis points. Panel (b) shows the realized spread, expressed in basis points. Panel (c) shows order book depth, expressed in thousands. Panel (d) shows volatility. Panel (e) shows currency volume, expressed in millions. Panel (f) shows share volume, expressed in thousands. In all plots, the left vertical break indicates June 1, 2009, the start of the 'tick size war'. The middle vertical break indicates July 6, 2009, the date of OSEs tick size reduction. The right vertical break indicates August 31, 2009, when tick sizes were harmonized across all exchanges.

Figure A.1: Time-series: Market quality OSE 2009

Panel A: 2008

Liq.measure	Market	May	1jun-5jul	6jul-31aug	Sep
Quoted (relative) Spread (%)	OSE	0.242	0.274	0.299	0.446
	Chi-X		0.415	0.629	0.666
Effective Spread (%)	OSE	0.087	0.097	0.102	0.145
	Chi-X		0.133	0.190	0.233
Realized Spread (%)	OSE	0.029	0.028	0.028	0.037
	Chi-X		0.025	0.033	0.053
Realized Volatility (%)	OSE	0.694	0.703	0.795	1.346
	Chi-X		0.361	0.415	0.661
Depth (Thousand NOK)	OSE	2625	2212	2477	1295
	Chi-X		341	303	246
Volume (Million NOK)	OSE	480	363	291	380
	Chi-X		1	4	6
Turnover (%)	OSE	1.23	0.89	0.84	1.33
	Chi-X		0.00	0.01	0.01

Panel B: 2010

Liq.measure	Market	May	1jun-5jul	6jul-31aug	Sep
Quoted (relative) Spread (%)	OSE	0.277	0.281	0.270	0.221
	Chi-X	0.499	0.487	0.438	0.352
Effective Spread (%)	OSE	0.085	0.084	0.084	0.078
	Chi-X	0.160	0.140	0.137	0.119
Realized Spread (%)	OSE	0.013	0.016	0.018	0.024
	Chi-X	0.035	0.037	0.043	0.036
Realized Volatility (%)	OSE	0.721	0.581	0.499	0.429
	Chi-X	0.502	0.448	0.390	0.313
Depth (Thousand NOK)	OSE	535	497	544	788
	Chi-X	238	202	166	182
Volume (Million NOK)	OSE	305	198	176	196
	Chi-X	14	12	12	9
Turnover (%)	OSE	1.07	0.83	0.60	0.64
	Chi-X	0.05	0.05	0.04	0.03

The table reports subperiod averages of a number of market quality measures at the OSE and Chi-X. Market quality measures: *Quoted (relative) spread*: The difference between best bid and best ask in the order book, divided by price. Averaged across all order books during a trading day. *Effective spread*: Difference between trade price and a pre-trade benchmark, relative to trade price. *Depth*: The total (NOK) amount outstanding at the best bid and ask. Volume: The total amount (in NOK) traded. All the variables are measured as averages over daily observations in the given time interval.

Table A.2: Liquidity measures, comparable subperiods, 2008 and 2010

B Distribution of stock prices and tick sizes at the OSE

Tick sizes for OSE listed stocks are determined by a step-function of prices — higher priced stocks have larger tick sizes (the tick size schedules are discussed in Section 1.4). To inform about the distribution of stock prices at the OSE, and therefore the range of possible tick sizes, Figure A.2 plots the distribution of (end-of-day) stock prices for our sample of stocks on the last trading day of May, 2009. The figure shows that most of our sampled stocks are priced below 150 NOK. The lowest stock price in our sample is 3.68 NOK while the highest stock price is 226.25 NOK.

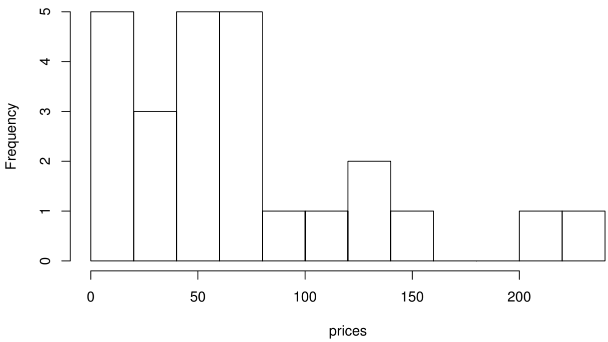

The histogram presents the distribution of stock prices at the Oslo Stock Exchange on the last trading day of May, 2009. The sample comprises all OBX index stocks. Stock prices are denominated in Norwegian Krone (NOK).

Figure A.2: Distribution of stock prices at the OSE (May, 2009)

C Regression discontinuity design

In Section 5, we use a difference-in-differences design to estimate the impact of the tick size war on market quality at the OSE and Chi-X. In this appendix, we instead use a so-called regression discontinuity design to estimate the impact of the tick size war on market quality. We also discuss why the regression discontinuity design is an improvement over the simple before-and-after specification presented in Appendix A.

C.1 Methodology

We propose a regression discontinuity methodology to identify the causal impact of the 'tick size war' for Nordic stocks on stock outcomes (e.g. trading quality, liquidity). While the 'tick size war' actually comprises two distinct events — the Chi-X tick size reduction for OSE listed stocks on June 1, 2009 and the OSE tick size reduction for OBX index stocks on July 6, 2009 — consider for now the evaluation of some arbitrary event implemented on date t^*, on the outcomes y_{it} for stock i on date t. One approach to assess the effect of event t^* on stock outcomes would be to use a before-and-after estimator:

$$y_{it} = \alpha + \gamma Event_t + \omega_{it}, \tag{1}$$

where

$$Event_t = \begin{cases} 1, & \text{if } t \geq t^* \\ 0, & \text{otherwise} \end{cases}$$

The before-and-after effect of interest is captured by the coefficient γ, while the error term ω_{it} represents all other determinants of the outcome. The coefficient γ is derived by computing the mean of y_{it} over *all* periods $t < t^*$, and subtracting it from the mean of y_{it} computed over *all* periods $t \geq t^*$. The coefficient γ, however, is unlikely to represent the causal impact of the events of the tick size war on outcomes y_{it}. The reason for this is that most of our outcome variables, such as stock liquidity and order book depth, are influenced by persistent trends that pre-date the tick size war (see the discussion in Section 3.4). Absent an adjustment for such pre existing trends, equation 1 will erroneously attribute the trends to the impact γ of the tick size war.

In this section, we approach the issue of pre-existing trends by focusing only on the variation in outcomes that occurs *exactly* on the date t^* of the event, in a regression discontinuity design. We conjecture that such local variation is unlikely to be correlated with other determinants of y_{it}, which may facilitate causal inference. We implement the regression discontinuity design with the following regression model:

$$y_{it} = \alpha + \underbrace{\beta^{Pre-trend}(t-t^*) + \beta^{Gradual}(t-t^*) \times Event_t}_{Added\,terms} + \beta^{Jump} Event_t + \varepsilon_{it} \tag{2}$$

where $(t - t^*)$ is an event-time counting variable, centered on the event date t^*. This variable is decreasingly negative for all dates leading up to t^*, and increasingly positive for all dates after t^*. Since event-time is centered on t^*, the coefficient β^{Jump} identifies a discrete change in y_{it} occurring exactly on the day of the event.[1] Similarly, $\beta^{Gradual}$ can be interpreted as the *per-day* impact of the event, identified by a change in the linear trend $(t - t^*)$ exactly on the day of the event. We estimate model (2) separately for the two markets, and for a variety of outcomes y_{it}.

[1] In contrast, in the 'traditional' before-and-after event-study methodology (equation 1), which does not include $(t - t^*)$ as a regressor, the coefficient on $Event_t$ captures the difference in mean outcomes before-and-after, where the means are computed over the entire 'before' and 'after' periods.

Our design shares both the strengths and the weaknesses of the regression discontinuity design. First, by focusing only on variation in outcomes close to $t = t^*$, the RD design gives the potential for causal inference, since such local variation is unlikely to be correlated with other determinants of y_{it}. Indeed, as long as there are no simultaneous shocks to y_{it} at t^*, the coefficients β^{Jump} and $\beta^{Phasein}$ capture the immediate and gradual causal effects of an event implemented at date t^*.

Consistent estimation of the coefficients β^{Jump} and $\beta^{Phasein}$, however, requires a strong assumption about the functional form of the relationship between $(t - t^*)$ and y_{it}. This assumption is needed because in order to estimate the effects that occur close to $t = t^*$, it is necessary to use data away from this point as well (Lee and Lemieux, 2010). Two main approaches are taken in the RD literature to estimate equation (2) when the functional form of $(t - t^*)$ is unknown. The first approach, which is widely preferred, is to restrict the sample size on either side of t^*, and estimate equation (2) with local linear regressions. If there is a concern that the regression function is not linear over the entire range of $(t - t^*)$, restricting the estimation range to values closer to the event date $t = t^*$ is likely to reduce biases in the RD estimates (Hahn, Todd, and van der Klaauw, 2001; Lee and Lemieux, 2010). The second approach, in contrast, uses all the available data and allows for a flexible relationship between y_{it} and $(t - t^*)$, by expanding equation (2) with polynomials in $(t - t^*)$.

We estimate equation (2) with local linear regressions, and restrict the amount of data we use before and after an event t^*. In order to do so, we make two definitions. First, we define the event dates t^* of interest. We wish to estimate the impact of the onset of the 'tick size war', on June 1, 2009, as well as OSEs tick size reduction on July 6, 2009. We label these events t_1^* and t_2^*, respectively. As equation (2) only allows us to center event-time around one event date at a time, we must estimate equation (2) separately for the events t_1^* and t_2^*.

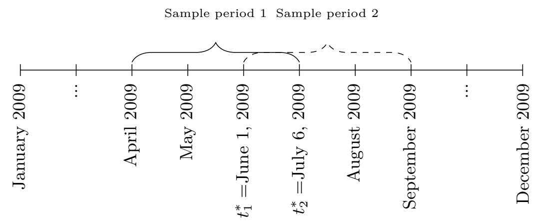

The figure illustrates how we define sample periods surrounding our two event dates. Our first event, t_1^*, is the beginning of the 'tick size war' on June 1, 2009. Our second event, t_2^*, is the Oslo Stock Exchange tick size reduction on July 6, 2009. First, surrounding the June 1, 2009, event, we restrict the sample period to April 1, 2009, to July 5, 2009. Second, surrounding the July 6, 2009 event, we restrict the sample period to June 1, 2009, to August 31, 2009. Solid curly braces span the sample period surrounding June 1, 2009. Dashed curly braces span the sample period surrounding July 6, 2009.

Figure A.3: Illustration: Sample restrictions

Second, we define sample periods separately for each of these events. In small event windows surrounding the events t_1^* and t_2^*, a linear approximation of the functional form of $(t - t^*)$ is likely to be appropriate. Figure A.3 illustrates how we restrict the sample periods surrounding both t_1^* and t_2^*. First, surrounding the June 1, 2009, event, we restrict the sample period to April 1, 2009, to July 5, 2009. Second, surrounding the July 6, 2009 event, we restrict the sample period to June 1, 2009, to August 31, 2009.[2] Figure A.4 provide compelling graphical

[2]The cutoff dates used to restrict the samples are far from arbitrary. For example, in restricting the sample

evidence that within these event windows, a linear functional form of $(t - t^*)$ indeed appears appropriate.

C.2 Results

Table A.3 shows results of this regression discontinuity analysis, allowing us to to estimate the impact of the events of the tick size war. In Panel A of the figure we document an unambiguously negative impact of the Chi-X tick size reduction on OSE market quality. For example, offering comparatively large tick sizes causes a *daily* exodus of OSE trading volume and order book depth by 6% and 2.7%, respectively, and, presumably as a consequence, a daily increase in effective spreads by 0.1 basis points — strongly suggesting a less liquid market.

In contrast, the impact of the Chi-X tick size reduction on Chi-X market quality is ambiguous. First, the coefficient estimates imply immediate improvements in Chi-X order book depth and trading volume, by respectively 14.5 and 72 per cent. These effects, however, appears to dissipate over time. Our estimates of $\beta^{Gradual}$ imply that out of the initial 14.5% (72%) improvement in depth (trading volume), 1.4 (5.9) percentage points dissipates per day. Moreover, there is evidence that spread measures of liquidity at Chi-X worsened during the tick size war.

Panel B of Table A.3 assesses the impact of the OSE tick size reduction on OSE market quality. Consistent with a voluminous empirical literature, we find a simultaneous and immediate decrease in both effective spreads and order book depth (−31.8%) following the OSE tick size reduction. At the same time, by reducing its tick sizes in line with its competitors, the OSE is able to abate the exodus of trading volume and order book depth spurred by the tick size war. This is indicated by highly significant and positive coefficient estimates of $\beta^{Gradual}$. In fact, the existing negative trend in OSE trading volume is fully reversed and becomes positive ($\beta^{Gradual} + \beta^{PreTrend} > 0$). Similarly, the existing negative trend in order book depth is nullified ($\beta^{Gradual} + \beta^{PreTrend} \approx 0$).

Trading at Chi-X appears to stabilize following the OSE retaliatory tick size reduction. For example, the volatility of prices at Chi-X declines significantly following OSEs tick size reduction. Moreover, the erratic trading volume at Chi-X appears to normalize — after a gradual decline in trading volume throughout the break-out phase, the trend tapers following OSEs tick size reduction (captured by $\beta^{Gradual}$).

---

period for the July 6, 2009 event, we end the sample on August 31, 2009, as this is the introduction date of FESE harmonized tick size schedules, and effectively the conclusion of the tick size war. Similarly, we begin that sample on June 1, 2009, so as to not sample data before the tick size war began. Doing so, however, means we have overlap between the two sample periods during June 2009. This is inevitable if we wish to estimate the impact of OSE's tick size reduction, on July 6.

Panel A: OSE

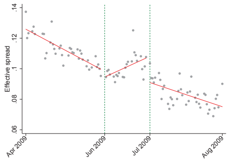

(a) Effective spread

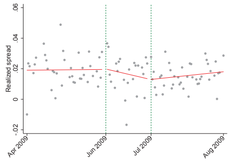

(b) Realized spread

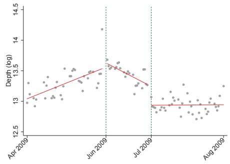

(c) Depth

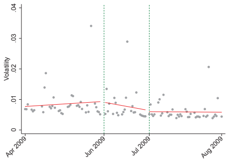

(d) Volatility

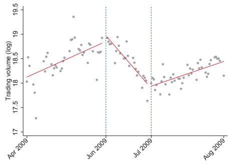

(e) Volume

Panel B: Chi-X

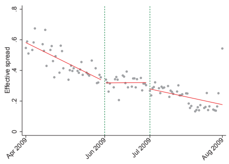

(f) Effective spread

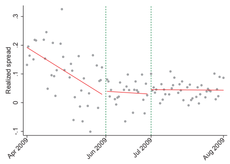

(g) Realized spread

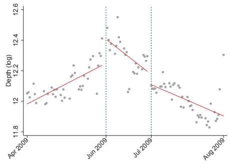

(h) Depth

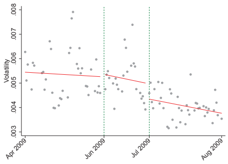

(i) Volatility

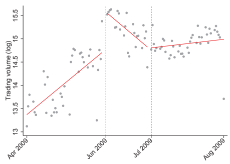

(j) Volume(NOK)

The figure presents time-series of market quality measures at the Oslo Stock Exchange (Panel A) and Chi-X (Panel B), in the period April 1, 2009, to August 31, 2009. All observations are daily cross-sectional averages, computed across all OBX listed stocks. Panel (a) shows the effective spread. Panel (b) shows the realized spread. Panel (c) shows order book depth, log-transformed. Panel (d) shows volatility. Panel (e) shows currency volume, log-transformed. In all plots, the left vertical break indicates June 1, 2009, the start of the 'tick size war'. The right vertical break indicates July 6, 2009, the date of OSEs tick size reduction. Linear regression lines (red) are fit separately within each event window. The regression lines correspond exactly with those generated in equation (2).

Figure A.4: Market quality OSE

Panel A: Chi-X tick size reduction ($t^* = $ June 1, 2009)

	Espread		*Rspread*		*Depth*		*Volatility*		*Volume*	
OSE										
$\beta^{Pre-trend}$	-0.001***	(-4.29)	0.000	(0.07)	0.012***	(8.08)	0.000	(1.38)	0.017***	(5.51)
$\beta^{Gradual}$	0.001***	(3.50)	-0.000	(-1.13)	-0.027***	(-7.78)	-0.000*	(-2.01)	-0.060***	(-12.95)
β^{Jump}	-0.003	(-1.16)	0.000	(0.01)	0.051	(1.24)	-0.001	(-0.47)	0.087	(1.42)
N	1612		1612		1612		1612		1612	
Adj. R^2	0.04		-0.00		0.03		-0.00		0.04	
CHI-X										
$\beta^{Pre-trend}$	-0.006***	(-5.10)	-0.004**	(-2.60)	0.006***	(4.31)	-0.000	(-0.58)	0.033***	(9.29)
$\beta^{Gradual}$	0.006***	(3.26)	0.004**	(2.75)	-0.014***	(-4.18)	-0.000	(-0.79)	-0.059***	(-7.52)
β^{Jump}	-0.006	(-0.20)	0.024	(0.50)	0.145***	(2.92)	0.000	(0.33)	0.723***	(6.98)
N	1412		1409		1550		1318		1497	
Adj. R^2	0.06		0.02		0.06		-0.00		0.15	

Panel B: OSE tick size reduction ($t^* = $ July 6, 2009)

	Espread		*Rspread*		*Depth*		*Volatility*		*Volume*	
OSE										
$\beta^{Pre-trend}$	0.001*	(1.81)	-0.000	(-1.42)	-0.015***	(-4.94)	-0.000*	(-1.90)	-0.043***	(-12.58)
$\beta^{Gradual}$	-0.001***	(-2.85)	0.000*	(1.72)	0.015***	(4.19)	0.000*	(1.78)	0.056***	(12.53)
β^{Jump}	-0.017***	(-4.89)	-0.000	(-0.05)	-0.318***	(-6.20)	-0.001	(-0.75)	0.009	(0.13)
N	1690		1690		1690		1689		1690	
Adj. R^2	0.05		0.00		0.11		0.00		0.05	
CHI-X										
$\beta^{Pre-trend}$	0.000	(0.02)	-0.000	(-0.43)	-0.009***	(-2.96)	-0.000	(-1.40)	-0.033***	(-3.50)
$\beta^{Gradual}$	-0.003	(-1.47)	0.000	(0.31)	0.004	(0.94)	0.000	(0.13)	0.038***	(3.26)
β^{Jump}	-0.037	(-1.70)	0.014	(0.68)	-0.070	(-1.29)	-0.001***	(-3.27)	0.025	(0.22)
N	1647		1647		1677		1590		1664	
Adj. R^2	0.04		-0.00		0.10		0.08		0.02	

The table presents regression discontinuity estimates of the impact of the tick size war (top panel), and OSEs tick size reduction (bottom panel), on market quality outcomes. *Espread* is the effective spread, in percentage points. *Rspread* is the realized spread, percentage points. *Depth* is the order book depth, log-transformed. *Volatility* is the realized volatility. *Volume* is the NOK trading volume, log-transformed. The regression specification is $y_{it} = \alpha + \beta^{Pre-trend}(t - t^*) + \beta^{Gradual}(t - t^*) \times Event_t + \beta^{Jump} Event_t + \varepsilon_{it}$, where $(t - t^*)$ is an event-time counting variable centered on the event date t^* (June 1, 2009 for top panel, July 6, 2009 for bottom panel). $Event_t$ is a dummy variable equal to 1 for all observations $t \geq t^*$. Surrounding the June 1 event, we restrict the sample period to April 1 to July 5. Surrounding the July 6 event, we restrict the sample period to June 1 to August 31. Standard errors are clustered at the stock-level. t-statistics in parentheses. ** $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$*

Table A.3: Market quality regressions

D Robustness tests: Difference-in-differences

In Section 5, we use a difference-in-differences specification to estimate the impact of the events of the tick size war on market quality at the OSE and Chi-X. In the difference-in-differences specification, we use a control group of stocks that are not directly affected by the tick size war to control for the influence of common confounding factors on our estimates of the effect of the tick size war. In this section, we explore the robustness of our benchmark difference-in-differences results to alternative specifications.

D.1 Benchmark difference-in-differences specification

Before describing our robustness tests, we begin by restating the benchmark difference-in-differences specification estimated in Section 5. In the benchmark model, we define two separate treatment groups. The first treatment group is OBX index stocks traded on the OSE. The second treatment group is OBX index stocks traded on Chi-X. Both groups were directly affected by the Chi-X tick size reduction for OSE listed stocks on June 1, 2009 and the OSE tick size reduction for OBX index stocks on July 6, 2009.

The control group is constructed in two steps. First, we construct a sample of 173 non-OBX index OSE stocks that were not traded on the multilateral trading facilities (MTFs) Chi-X, Turquoise, or BATS throughout the calendar year 2009. Since these stocks were not traded on the three MTFs, they were not directly affected by the MTF tick size reductions during June 2009. Moreover, since these stocks did not belong to the OBX index, they were not directly affected by the OSE tick size reduction on July 6, 2009. Second, we reduce the initial control sample of 173 stocks to the 25 most-traded stocks based on overall trading volume in the month of May, 2009, in order to provide a more comparable control group to our highly liquid treatment group.

We implement the difference-in-differences design with the following regression model:

$$y_{it} = \alpha_i + \alpha_t + \tau Treatment_{it} + \omega_{it}, \tag{3}$$

where $Treatment_{it} = 1$ for stock i that belongs to the treatment group on date $t \geq t^*$ and zero otherwise; α_i are stock-level fixed effects; and α_t are date-level fixed effects. The inclusion of stock and date fixed effects in equation 3 controls for constant differences in y_{it} between treatment and control sample stocks and ensures that the effect of $Treatment_{it}$ on y_{it} is measured net of the time trend in the control sample.

Equation 3 is estimated separately for the two events of interest — the Chi-X tick size reduction on June 1, 2009 and the OSE tick size reduction on July 6, 2009. As in the main text, surrounding the June 1 event, we restrict the sample period to April 1 to July 5. Surrounding the July 6 event, we use a sample period from June 1 to August 31.

Table 4 in the main text presented estimates from the benchmark difference-in-differences model. The table suggests that both our spread measures of liquidity at the OSE deteriorated as a result of the June 1, 2009, Chi-X tick size reduction while Chi-X depth and trading volume increased. The table also shows that order book depths at both the OSE and Chi-X declined considerably following the July 6, 2009, OSE tick size reduction.

D.2 Robustness test: Alternative control samples

Our first robustness test is to estimate the benchmark difference-in-differences design using two alternative control group specifications. The first alternative control group, which we label Control group 1, comprises all 173 non-OBX index OSE stocks that were not traded on the multilateral trading facilities (MTFs) Chi-X, Turquoise, or BATS throughout the calendar year 2009. Recall that the benchmark control group comprises the 25 most-traded stocks from Control group 1. The second alternative control group, which we label Control group 2, retains from Control group 1 only stocks with positive trading volume at least 200 out of the 251 trading days during the calendar year 2009. This requirement excludes the least liquid stocks from the control sample and potentially creates a better comparison group for our liquid OBX index treatment group. Control group 2 holds 81 stocks.

To facilitate a comparison between our three control group specifications, in Table A.4, we present summary statistics from both the benchmark control group specification, Control group 1, and Control group 2. The table illustrates that stocks in the benchmark control group are the most liquid while stocks in Control group 1 are the least liquid. This is no surprise, as both the benchmark control group and Control group 2 are derived from Control group 1 conditional on a parameter of stock liquidity.

	μ	σ	Min.	Median	Max.	N
Benchmark control group						
Effective spread (%)	0.813	0.579	0.154	0.706	7.113	2162
Realized spread (%)	0.440	0.655	-5.128	0.310	8.333	2162
Depth (thousand NOK)	177.087	159.135	17.445	126.398	1677.866	2515
Volatility (%)	1.037	0.845	0.052	0.874	15.012	1948
Trading volume (thousand NOK)	4753.817	19930.033	5.400	1349.027	650334.994	2408
Control group 1						
Effective spread (%)	1.312	1.042	0.000	1.021	9.285	7652
Realized spread (%)	0.753	1.320	-17.407	0.470	16.988	7605
Depth (thousand NOK)	193.372	2006.586	3.633	83.152	77185.746	14678
Volatility (%)	1.188	1.190	0.026	0.921	36.017	5498
Trading volume (thousand NOK)	1766.009	15014.750	1.000	196.785	882447.457	11565
Control group 2						
Effective spread (%)	1.151	0.825	0.000	0.947	7.637	6068
Realized spread (%)	0.650	1.009	-5.128	0.441	13.967	6036
Depth (thousand NOK)	120.527	119.086	6.060	84.583	1677.866	8100
Volatility (%)	1.112	0.995	0.052	0.896	36.017	4652
Trading volume (thousand NOK)	1952.345	11910.515	1.600	327.423	650334.994	7586

The table presents summary statistics from our three difference-in-differences control group specifications. The baseline control group consists of non-OBX index OSE stocks that were not traded on the multilateral trading facilities (MTFs) Chi-X, Turquoise, or BATS throughout the calendar year 2009. In Control group 1, we further restrict the control sample to stocks with more than 200 trading days throughout 2009. In Control group 2, we restrict the sample to the 25 most traded stocks in the baseline control group based on total trading volume during May 2009. Summary statistics are computed using observations from January to May, 2009. The table lists means (μ), standard deviations (σ), minimum (Min.) and maximum values (Max.), medians, and number of observations (N).

Table A.4: Summary statistics control sample

Estimates from the difference-in-differences model using Control group 1 are presented in Table A.5, labeled as specification 1. In the top panel of Table A.5, we assess the impact of the Chi-X tick size reduction ($t_1^* =$ June 1, 2009) on the quality of trading at the OSE and Chi-X. The table shows that stock liquidity at the OSE deteriorates as a result of the Chi-X tick size reduction. For example, effective spreads increase by 0.88 percentage points for OSE listed stocks directly affected by the tick size reduction relative to a control group of OSE listed stocks not affected by the tick size reduction. We find only weak evidence that Chi-X market quality increased, despite capturing market shares from the OSE (see section 4). In particular, effective and realized spreads decrease and depth increases but these effects are all statistically insignificant.

In the bottom panel of Table A.5, we evaluate the impact of the OSE tick size reduction (t_2^*=July 6, 2009) on stock market quality. The OSE tick size reduction causes a considerable reduction in order book depth at both the OSE (-45%) and Chi-X (-22%) — both effects measured relative to OSE listed stocks with no tick size change. At the same time, we find no impact of the OSE tick size reduction on spread measures of liquidity or volatility at neither the OSE nor Chi-X.

Estimates from the difference-in-differences model using Control group 2 are presented in Table A.5, labeled as specification 2. The results in Table A.5 support our previous findings that the Chi-X tick size reduction on June 1, 2009 adversely affected stock liquidity at the OSE, and that the OSE tick size reduction on July 6, 2009 reduced order book depths at both the OSE and Chi-X. In addition, the table provides weakly statistically significant evidence that stock liquidity at Chi-X (measured by effected spreads) improved during the tick size war.

Panel A: Chi-X tick size reduction ($t^* =$ June 1, 2009)

	Effective spread			*Realized spread*			*Depth*			*Volatility*		
	1	2	3	1	2	3	1	2	3	1	2	3
τ	0.088***	0.077***	0.088*	0.047	0.051*	0.041	0.035	0.027	0.079	-0.001	-0.000	-0.003
	(3.21)	(2.76)	(1.97)	(1.35)	(1.69)	(1.19)	(0.66)	(0.52)	(0.90)	(-0.56)	(-0.27)	(-1.45)
N	6717	5729	998	6698	5716	999	10615	6620	1019	5358	4903	981
Adj. R^2	0.61	0.60	0.63	0.24	0.25	0.18	0.76	0.84	0.88	0.14	0.09	0.05
CHI												
τ	-0.064	-0.075*	0.013	-0.027	-0.021	0.014	0.082	0.074	0.148*	-0.000	0.000	-0.001
	(-1.61)	(-1.86)	(0.27)	(-0.56)	(-0.47)	(0.23)	(1.47)	(1.36)	(1.69)	(-0.38)	(0.38)	(-0.75)
N	6524	5536	955	6502	5520	956	10564	6569	999	5066	4611	915
Adj. R^2	0.56	0.52	0.50	0.23	0.22	0.14	0.60	0.68	0.72	0.26	0.24	0.07

Panel B: OSE tick size reduction ($t^* =$ July 6, 2009)

	Effective spread			*Realized spread*			*Depth*			*Volatility*		
	1	2	3	1	2	3	1	2	3	1	2	3
τ	0.021	0.029	-0.042	-0.177	-0.216	-0.045	-0.446***	-0.458***	-0.354***	-0.002	-0.002	0.001
	(0.62)	(1.03)	(-1.01)	(-1.11)	(-1.19)	(-0.79)	(-5.79)	(-5.58)	(-4.22)	(-1.02)	(-0.83)	(0.53)
N	6566	5807	950	6545	5788	949	11023	6935	1007	5230	4887	903
Adj. R^2	0.61	0.57	0.73	0.43	0.48	0.38	0.71	0.77	0.86	0.35	0.39	0.04
CHI												
τ	-0.053	-0.045	-0.067	-0.165	-0.203	-0.032	-0.221***	-0.233***	-0.091	-0.001	-0.001	-0.001
	(-1.33)	(-1.32)	(-1.50)	(-1.03)	(-1.11)	(-0.52)	(-2.84)	(-2.81)	(-1.01)	(-0.79)	(-0.58)	(-0.67)
N	6523	5764	937	6502	5745	936	11010	6922	1007	5131	4788	864
Adj. R^2	0.58	0.52	0.65	0.42	0.47	0.31	0.59	0.59	0.65	0.47	0.54	0.16

The table presents estimates of τ from the difference-in-differences specification applied separately to the Chi-X tick size reduction ($t^* =$ June 1, 2009) and the OSE tick size reduction ($t^* =$ July 6, 2009). Surrounding the June 1, 2009, event, we restrict the sample period to April 1, 2009, to July 5, 2009. Surrounding the July 6, 2009 event, we restrict the sample period to June 1, 2009, to August 31, 2009. The difference-in-differences is estimated for three different robustness specifications, labeled 1, 2 and 3. In specification 1, the control group comprises all 173 non-OBX index OSE stocks that were not traded on Chi-X, Turquoise, or BATS Europe throughout the calendar year 2009. In specification 2, we further restrict the control sample to only comprise stocks with 200 or more trading days during the calendar year 2009. In specification 3, we restrict the sample period to 10 trading days before and after each of the two events (June 1, 2009 and July 6, 2009), using the same control sample as in Section 5. The difference-in-differences regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Treatment_{it} + \omega_{it}$. $Treatment_t$ is a dummy variable equal to 1 for all treatment group observations on dates $t \geq t^*$. The difference-in-differences specification is estimated separately for two treatment groups. The first treatment group is OBX index stocks traded on the OSE. The second treatment group is OBX index stocks traded on Chi-X. *Espread* is the effective spreads, in percentage points. *Rspread* is the realized spreads, in percentage points. *Depth* is order book depth, transformed with the natural logarithm. *Volatility* is measured in percentage points. Standard errors are clustered at the stock-level.

Table A.5: Difference-in-differences robustness tests

D.3 Robustness test: Shorter sample period

Finally, we return to the benchmark control group specification but shorten the sample period surrounding our two event dates (June 1, 2009 and July 6, 2009) to reduce the potential for confounding factors influencing our estimates. Specifically, we restrict the sample period to ten trading days before and after each of the event dates. Estimates from this robustness test are presented in Table A.5, labeled as specification 3. Shortening the sample period increases the noise in our estimates but our main empirical conclusions remain the same. In particular, we find that the Chi-X tick size reduction increased effective spreads at the OSE, and that the OSE tick size reduction reduced order book depth at the OSE by appreciable amounts.

References

Hahn, J., P. Todd, and W. van der Klaauw (2001): "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, 69, 201–209.

Lee, D. S. and T. Lemieux (2010): "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 48, 281–355.

Chapter 3

Tick sizes in illiquid order books

Tick sizes in illiquid order books

Tom Grimstvedt Meling*

June 21, 2017

Abstract

I assess the causal impact of increasing the tick size on stock liquidity and trading volume in illiquid stocks. Using a regression discontinuity design at the Oslo Stock Exchange, I find that increasing the tick size has no impact on the transaction costs, order book depths, or trading volumes of illiquid stocks. These findings contradict recent theoretical predictions in the market microstructure literature as well as proposals by lawmakers in the United States to increase the tick size for illiquid stocks.

*University of Bergen, Department of Economics. Tom.Meling@econ.uib.no. This paper has benefited from discussions with Hans K. Hvide, Katrine V. Løken, Bjørn Sandvik, and Eirik A. Strømland. Seminar participants at the University of Bergen have provided valuable insights. I am very grateful to Bernt Arne Ødegaard for sharing data with me. All errors are my own.

Introduction

Stock exchanges fine-tune their market designs to improve liquidity. A much-used strategy over the last two decades has been to reduce the tick size — the smallest price increment on an exchange.[1] However, the impact of tick size reductions on stock liquidity is uncertain. On the one hand, a smaller tick size can enhance price competition among investors and lead to narrower bid-ask spreads. On the other hand, a smaller tick size makes it easier to undercut other investors' limit orders, which can discourage investors from providing liquidity with limit orders. This ambiguity has created strong demand among policy makers for evidence on the impact of tick sizes on stock liquidity, in particular for illiquid stocks.[2]

The purpose of this paper is to assess the causal impact of tick sizes on stock liquidity and trading volume for both liquid and illiquid stocks. Buti et al. (2015) show theoretically that tick size reductions can decrease liquidity in illiquid stocks but increase liquidity in liquid stocks. The mechanism behind their result is that tick size reductions for liquid stocks enhance price competition, resulting in narrower bid-ask spreads and increased aggregate depth (though depth at the best bid-ask declines). However, as traders switch from market orders to limit orders, total trading volume declines. For illiquid stocks, in contrast, Buti et al. (2015) show that the costs of discouraging liquidity supply dominate the benefits of enhancing price competition, such that a reduction in the tick size reduces order book depth and widens the bid-ask spread, while total trading volume increases.

A regression discontinuity design at the Oslo Stock Exchange (OSE) allows for clean identification of the effect of tick sizes on stock liquidity and trading volume. I exploit that tick sizes at the OSE are determined as a function of the stock price — higher priced stocks have larger tick sizes. Comparing stocks that are priced marginally above tick size price thresholds to stocks that are priced marginally below the price thresholds in a regression discontinuity design allows for causal inference.

[1] For example, tick sizes in the United States have gradually declined over the past decades. The American Stock Exchange (AMEX) reduced its tick size for selected stocks to \$1/16 in 1992, and further applied this tick size to all AMEX stocks in 1997. Also in 1997, the New York Stock Exchange and NASDAQ implemented \$1/16 tick sizes. Decimal pricing was phased in from 2000, and was fully implemented by 2001.

[2] As a means to learn more about the effects of tick sizes on the liquidity in small and illiquid securities, policy makers in the United States have recently initiated a large-scale experimental program that has increased the tick size for 1200 randomly chosen small capitalization securities. The 'Tick Size Pilot Program' officially commenced in late 2016 and will last for a two-year period.

I use the regression discontinuity design to explore the causal effect of tick sizes on the liquidity in liquid stocks. To this end, I explore a long sample period ($2008 - 2011$) with exogenous variation in the tick size for the most liquid stocks at the OSE — the 25 stocks in the OBX index. I find that increasing the tick size for this population of liquid stocks leads to wider spreads and increased order book depth at the best bid and ask. Moreover, the regression discontinuity design shows a weak and potentially time-varying positive impact of increasing the tick size on trading volume. These results are broadly consistent with the theoretical predictions in Buti et al. (2015) for liquid order books.

To explore the effects of tick size changes for illiquid stocks, I apply the regression discontinuity design to a sample comprising a large number of both liquid and illiquid stocks at the OSE (all non-OBX index stocks). For this population of stocks, there are more than 2300 exogenous tick size changes distributed across 158 unique stocks in the period $2008 - 2011$, allowing for precise estimation of both average treatment effects and effect heterogeneity. I find that the average causal effect of increasing the tick size for the combined sample of liquid and illiquid stocks is to widen bid-ask spreads and to increase order book depth. However, the average effect is mostly accounted for by the most liquid stocks (top 40% of the liquidity distribution), whose liquidity responds heavily to tick size changes. In contrast, I find no impact of tick size changes on spread measures of liquidity, order book depth, volatility, or trading volume for stocks in the bottom 60% of the liquidity distribution.

This paper connects to several academic debates. First, my results connect to the already voluminous empirical literature on the impact of tick sizes on measures of stock liquidity (for a recent survey of the literature, see SEC 2012). The existing empirical literature has mostly focused on one-off tick size reforms where identification is difficult.[3] Similar to Buti et al. (2015), I exploit exogenous variation in tick sizes in a regression discontinuity design for causal inference. In line with Buti et al. (2015), I find that the average effect of increasing the tick size is to widen spreads and to increase order book depth.

Second, I contribute to the emerging empirical literature which explores whether tick

[3]Much of the existing literature is based on before-and-after variation in tick sizes surrounding regulatory reforms, which does not allow for a separation of the effect of tick sizes from confounding trends (e.g., Goldstein and Kavajecz 2000, Ronen and Weaver 2001). Some papers attempt to adjust for confounding trends by estimating the effects of tick size reforms net of the trend in a control sample of unaffected stocks (e.g., Bacidore et al. 2003, Chakravarty et al. 2004). This approach captures the causal effect of tick sizes only under the strict assumption that reform stocks and control stocks follow the same trends in the absence of tick size reform.

sizes affect liquid and illiquid stocks differently. Buti et al. (2015) build a theoretical model which predicts opposite effects of tick size changes for liquid and illiquid stocks, and test their predictions using data from the London Stock Exchange, NYSE, and Nasdaq. However, as the authors themselves point out, their data are ill-suited for testing predictions related to illiquid order books as most of their sampled stocks are, in fact, liquid.[4] In contrast, the Oslo Stock Exchange comprises a wide range of both liquid and illiquid stocks, which allows me to test the causal impact of tick size changes in both liquid and illiquid trading environments. Doing so, I find that the quality of trading in liquid stocks responds heavily to tick size changes while the quality of trading in illiquid stocks is unaffected by tick size changes.[5]

Finally, my research can provide guidance to policy makers in the United States who are currently considering tick sizes as a tool to improve the quality of trading in illiquid securities (see footnote 2). My causal estimates suggest that other market design tools than tick sizes are needed if the object is to improve the quality of trading in illiquid stocks.

The paper proceeds as follows. Section 1 provides institutional background on the determination of tick sizes at the Oslo Stock Exchange; Section 2 describes the data; Section 3 estimates a benchmark before-and-after event study specification; Section 4 describes the empirical identification strategy; Section 5 presents the main results; and Section 6 discusses the results and concludes.

[4]Buti et al. (2015) test their theoretical predictions using three data samples and two different empirical designs; a regression discontinuity design to exploit a price-based tick size for liquid securities at the London Stock Exchange; a regression discontinuity design to exploit that the tick size for stocks in the United States increases from $0.0001 to $0.01 as they cross the $1 price threshold; and a Fama-MacBeth approach to explore how changes in the relative tick size affect a sample of 180 NYSE and Nasdaq stocks. Among these data samples, only U.S. securities surrounding the $1 price threshold can plausibly be defined as illiquid. Nevertheless, Buti et al. (2015) use their estimates from the low-priced U.S. sample to shed light on theoretical predictions concerning liquid stocks. A potential explanation for why the authors choose not to explore in greater detail how the effect of crossing the $1 price threshold depends on initial stock liquidity, is that their sample of low-priced U.S. stocks only comprises 20 unique securities.

[5]In other empirical work, O'Hara et al. (2015) explore whether changes to the relative tick size affect stocks in a one-tick environment (the bid-ask spread is equals the tick size) and stocks in multi-tick environments differently. They show that in the one-tick environment, an increase in the relative tick size leads to more trading volume and increased order book depth. In contrast, in the multi-tick environment an increase in the relative tick size leads to less trading volume and less order book depth. My results connect to O'Hara et al. (2015) since my classification of liquid and illiquid stocks captures a similar separation between one-tick and multi-tick trading environments. In particular, the most liquid stocks in my sample tend to trade in (or close to) one-tick environments while the least liquid stocks tend to trade in multi-tick environments. Unlike O'Hara et al. (2015), I find no effect of tick size changes in multi-tick environments but a strong effect of tick size changes in one-tick environments.

1 Institutional background

This section gives an overview of the market design and institutional setting of the Oslo Stock Exchange before it describes in detail how tick sizes are determined at the Oslo Stock Exchange.

1.1 Overview: The Oslo Stock Exchange

The OSE operates a fully electronic limit order book, and has done so since January 1999. The OSE order book allows conventional limit orders, market orders, iceberg orders and various other common order types. Order placements at the OSE follow price-time priority — orders are first sorted by their price and then, in case of equality, by the time of their arrival.[6] The trading day at the OSE comprises three separate trading sessions: an opening call period, a continuous trading period, and a closing call period. In late 2012, the continuous trading session was shortened from 09:00 – 17:20 to 09:00 – 16:20. Call auctions may be initiated during continuous trading if triggered by price monitoring or to restart trading after a trading halt. Meling (2016) provides details on the market transparency at the OSE.

Competing stock exchanges offer trading in some, but not all, of the $200 - 300$ stocks listed at the OSE. In 2008, competing stock exchanges offered trading only in the largest and most liquid stocks at the OSE, before gradually expanding their selection of tradable stocks. For example, Chi-X, a so-called multilateral trading facility (MTF), initially offered trading in only the five largest OSE stocks (Norsk Hydro ASA, Renewable Energy Corp. A/S, StatoilHydro ASA, Telenor ASA, and Yara International ASA). At the time of writing in 2016, Chi-X offers trading in more than 50 OSE products. Likewise, Turquoise initially opened trading in 28 OSE stocks in 2008 but has since greatly expanded its selection to include more than 150 OSE products. For more details on the exchange competition for order flow in OSE listed products, see Meling and Ødegaard (2016).

[6] After the sample period I study, the OSE has adopted a price-visibility-time priority scheme where for price equality displayed orders are given preference over hidden orders. Traders also have the option to preferentially trade with themselves before trading with other traders. Such orders execute according to price-counterparty-visibility-time.

1.2 Tick sizes at the Oslo Stock Exchange

Tick sizes at the OSE are determined as a function of stock prices — stocks with higher prices have larger tick sizes. When prices cross a pre-specified price threshold from below (above) the tick size increases (decreases) instantly and automatically. I refer to the combined set of stock price thresholds that determine tick sizes as a 'tick size schedule.'

Over the last decade, there have been several changes to the tick size schedules at the OSE. Table 1 summarizes all the tick size schedules used by the OSE in the period 2003 – 2012. From June 2003, all stocks at the OSE shared the same 'four-step' tick size schedule with price thresholds at 10NOK, 50NOK, 150NOK, and 1000NOK. In September 2006, the OSE introduced separate tick size schedules for its large-cap stocks and small-cap stocks. Stocks listed on the OBX index, which contains the 25 most traded stocks at OSE, are defined by the OSE as 'large caps.'[7]

The tick size schedules introduced in September 2006 were maintained until the Summer of 2009, when a 'tick size war' erupted between the OSE and several competing stock exchanges (the events of this tick size war are described in detail by Meling and Ødegaard 2016). Beginning on June 1, 2009, Chi-X significantly reduced the tick size for its selection of OSE listed stocks, quickly followed by Turquoise (June 8) and BATS Europe (June 15). On July 6, 2009, the OSE responded by reducing the tick size for all OBX index to a flat 0.01NOK. On August 31, 2009, all stock exchanges agreed on and implemented a shared pan-European tick size schedule for OBX index stocks, mandating much smaller tick sizes than before the tick size war.

2 Data

In this section, I describe the data sources used in this study and define measures of stock liquidity. Finally, I provide summary statistics from the data sample.

[7]The OBX index is aimed to be a highly liquid composition of shares that reflects the Oslo Stock Exchange investment universe. The stock composition of the OBX is revised twice a year (end of June and December). Stocks are selected for the OBX list based on cumulative trading volume in the six months leading up to a new OBX composition. For trading at the OSE, the OBX shares tend to have different rules than the other shares listed at the OSE (see for example Meling 2016).

2.1 Data sources

I employ two datasets to inform about the impact of changing the tick size on stock market quality at the Oslo Stock Exchange. First, I collect daily frequency data on all common stock at the Oslo Stock Exchange from *Børsprosjektet* at the Norwegian School of Economics (similar to CRSP). The data covers the period January 2003 - December 2011. This dataset holds information on opening and closing prices, daily price dispersion (highest and lowest prices), measures of trading volume (in NOK and in shares), end-of-day bids and asks, and OBX and OSEBX index constituency indicators. I generate tick sizes from these data on a daily level based on information on end-of-day prices and the prevailing tick size schedule for a given stock (Table 1).

Second, to explore how tick sizes affect measures of stock liquidity and trading costs, I use the ThomsonReuters Tick History (TRTH) Database. The TRTH database contains trade-and-quote data for OSE listed stocks across all European equity market places, and is available in the time period 2008 – 2011. For lit exchanges (where the limit order book is displayed), the TRTH provides information on the ten best levels of the bid and ask side of the limit order book. The ThomsonReuters data also includes information on over-the-counter trading of OSE shares, by including trades reported by Markit BOAT (a MiFID-compliant trade reporting facility).

2.2 Sample selection

In the main empirical analysis (Section 5), I place three restrictions on the data. First, I exclude from the overall data sample (January 2008 – December 2011) observations in the time period June 2009 — August 2009, a highly disruptive period where competing stock exchanges challenged OSE market shares by reducing tick sizes for OSE listed stocks (see Meling and Ødegaard 2016).

Second, I restrict the sample based on stock prices. While the OSE tick size schedules provide exogenous variation in tick sizes up to the 1000NOK price threshold, there is only sufficient variation around the lower-priced thresholds. To illustrate this point, Figure 1 plots the frequency of observations at each stock price level for both non-OBX and OBX index stocks. In order to have sufficient data surrounding each of the tick size thresholds, I remove from both the OBX and non-OBX samples all stocks whose price exceeds 200NOK at any point in time throughout the sample period $2008-2011$. Furthermore, the tick size price thresholds for low-priced OBX stocks are closely spaced, especially in the time period September $2009-2011$, which reduces the amount of data available around each threshold (see Table 1). To circumvent this issue, I remove from the OBX sample stocks whose price at any point in time during $2008-2011$ falls below 5NOK.

Notice, however, that the sample restrictions described above do not apply to the benchmark before-and-after analysis in Section 3. In the before-and-after analysis, I use data from the time period June 2009 — August 2009 and place no price-based restrictions on the data.

2.3 Variable construction

I use the ThomsonReuters Tick History database to compute a variety of stock liquidity measures. To capture the transaction cost dimension of stock liquidity, I compute two spread measures of liquidity. First, the relative spread is defined as the difference between the current best bid and ask divided by the quote midpoint. The relative spread is updated whenever the limit order book is updated, and is calculated as the average of these estimates throughout the trading day.

Second, the realized spread captures the gross revenue to liquidity suppliers after accounting for adverse price movements following a trade. The 5-minute realized spread for transaction j in stock i is given by $q_{ji}(p_{ji} - m_{i,j+5\text{min}})/m_{ji}$, where q_{ji} is an indicator variable that equals $+1$ for buyer-initiated trades and -1 for seller-initiated trades; p_{ji} is the trade price; and $m_{i,j+5\text{min}}$ is the quote midpoint 5 minutes after the j'th trade. To determine whether an order is buyer or seller initiated, the transaction price is compared to the previous quote midpoint — if the price is above (below) the midpoint it is classified as a buy (sell). The daily realized spread is computed as the average across all transactions during the trading day.

The depth dimension of stock liquidity is captured by calculating the sum of pending trading interest at the best bid and ask prices, measured in monetary terms (NOK). My measure of order book depth is updated whenever the limit order book is updated, and averaged across all order book states throughout the trading day. To proxy for the noise in the price process, I estimate realized volatility as the second (uncentered) sample moment of the within-day 10-minute stock returns.

Since the liquidity measures described above are based on within-day data while tick sizes in my setting are based on end-of-day stock prices, regressions of liquidity outcomes on tick sizes may be affected by measurement error. For example, a stock may cross a tick size price threshold during the trading day and cross back below the price threshold before the close. The end-of-day tick size would not reflect these price crossings but the liquidity measures might. Such measurement error, however, should only serve to attenuate the regression discontinuity estimates.

2.4 Summary statistics

Table 2 presents summary statistics of the trading in both small-cap and large-cap stocks at the OSE. All summary statistics are based on the Reuters order-level data from the time period $2008 - 2011$. The table shows that trading in small-cap stocks at the OSE differ from large-cap trading in several ways. First, there are considerable differences in stock liquidity, measured both in transaction costs and in order book depth. For example, the relative spread is (on average) 369.42 basis points (bps.) in small-caps and only 29.47 bps. in large-caps. Similarly, the realized spread is 59.74 bps. and 2.21 bps for small and large-caps, respectively. Large-cap order books are more than twice as deep as small-cap order books, and the average trading volume in large-cap stocks (155 million NOK) is more than 30 times larger than the trading volume in small-cap stocks (4.77 million NOK). Perhaps as a result of the greater liquidity, price volatility in large-caps is considerably smaller than in small-caps.

Second, tick sizes, both in absolute terms and in relative ($ticksize/price$) terms, differ between liquid and illiquid stocks at the OSE. In particular, tick sizes for small-caps are larger than for large-caps even though small-cap stock prices are lower. This is because the large-cap tick size schedules mandate smaller tick sizes for any given stock price. As a consequence, the relative tick size is five times larger for small-caps than for large-caps. At the same time, the tick size appears to be a less binding constraint for small-caps than for large-caps. For example, the 'ticks-per-quoted-spread', a common measure of how binding the tick size is, averages 3.69 for the large-cap sample and 10.44 for small-caps. Thus, the likelihood of the tick size being a binding constraint on the bid-ask spread differs considerably between the large-cap and small-cap samples.

3 Benchmark methodology: Before-and-after

To provide a benchmark for my later regression discontinuity estimates, and to replicate the methodology used in much of the existing empirical literature, I begin my empirical analysis by estimating the impact of tick size changes on stock market quality using a simple before-and-after specification. On July 6, 2009, the OSE unilaterally reduced the tick size for the 25 stocks in the OBX index to 0.01NOK. Before this date, tick sizes for OBX index stocks were determined by individual stock prices, and the stock price mandated tick sizes were typically much larger than 0.01NOK (see Table 1 for the full tick size schedules).

I estimate the impact of the July 6, 2009 tick size reduction using the standard before-and-after estimator:

$$y_{it} = \alpha + \beta Post_t + \epsilon_{it}, \tag{1}$$

where $Post_t = 1$ for observations after the event date July 6, 2009. Consequently, the regression coefficient β captures the difference-in-means in y_{it} before and after the event date, which is typically interpreted as a measure of the effect of the tick size change on y_{it}. I estimate equation 1 using a short sample period surrounding the event date — ten trading days before and ten trading days after the event date — to minimize the influence of confounding factors on my estimate of β.

Table 3 presents estimates from the before-and-after specification. The table shows that, in line with the existing empirical research, the OSE tick size reduction leads to tighter relative spreads (-10%, $t-stat = 2.27$) and shallower order books (-42%, $t-stat = 9.36$). Moreover, the before-and-after exercise reveals that reducing the tick size leads to less trading activity, captured by a 12% reduction in NOK trading volume ($t - stat = 2.14$). I find no impact of the tick size reduction on realized spreads or volatility.

4 Methodology

The purpose of this section is to devise an empirical methodology which can estimate the causal relationship between tick size changes and measures of stock liquidity and trading volume. In Section 3, I used a before-and-after estimator to assess the effect of a tick size reduction on stock outcomes:

$$y_{it} = \alpha + \beta Post_t + \epsilon_{it}, \tag{2}$$

where

$$Post_t = \begin{cases} 1, & \text{if } t \geq t^* \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

and t^* = July 6, 2009 — the date of a tick size reduction for the most liquid stocks at the OSE. The before-and-after effect of interest is captured by the coefficient β, while the error term ϵ_{it} captures all other determinants of the outcome. The coefficient β is derived by computing the mean of y_{it} over all periods $t < t^*$, and subtracting it from the mean of y_{it} computed over *all* periods $t \geq t^*$. In Section 3, the estimates of β suggested that reducing the tick size for liquid stocks results in a reduction in both spread measures of liquidity and order book depth, and a reduction in trading volume.

The coefficient β, however, is unlikely to capture a causal relationship between tick sizes and outcomes y_{it}. The reason for this is that before-and-after estimators, in general, are notoriously susceptible to the influence of pre-existing trends and seasonal effects. The setting surrounding the July 6, 2009 tick size reduction at the OSE is no different — for example, Meling and Ødegaard (2016) point out that stock liquidity at the OSE was improving throughout the calendar year 2009 for reasons unrelated to tick size reductions, and that trading behavior at the OSE tends to be different during the Summer months even in the absence of tick size changes. As a consequence, $Post_t$ may be correlated with omitted variables that are themselves correlated with y_{it} — leading to a biased estimate of β.

The price-based tick size determination at the Oslo Stock Exchange provides a useful source of exogenous variation to overcome this endogeneity problem. Stocks that are priced marginally above a tick size price threshold are assigned to a different tick size than stocks that are priced marginally below a tick size threshold. If traders cannot (or will not) strategically manipulate prices in order to induce tick size changes, it is essentially random whether a stock is priced marginally above or marginally below a tick size threshold.[8]

The so-called regression discontinuity (RD) design can be used to exploit such quasi-

[8] Such strategic pricing behavior would most likely result in a discontinuous change in the density of price observations at the tick size price thresholds (McCrary 2008). Reassuringly, however, Figure 1 indicates that there is no excess density (or bunching) at the price levels where the tick sizes increase, suggesting an absence of price manipulation which could invalidate the empirical design.

random variation. The RD design relates discontinuities in outcomes at some 'treatment' threshold to discontinuities in the probability of treatment at the same point (see Lee and Lemieux 2010 for a survey). In the context of tick sizes at the Oslo Stock Exchange, the RD design relates discontinuities in the tick size (panel a, Figure 2) to discontinuities in outcomes at the same price levels (panel b, Figure 2). The basic idea is that stocks that are priced, for example, 49NOK are likely to provide an adequate control group for stocks that are priced 50NOK. In such a setting, differences in outcomes between stocks priced marginally above and marginally below a price threshold can be attributed to the difference in tick size that the two stocks experience.

To implement the RD approach in my empirical setting, with a discrete treatment variable of interest (as opposed to binary) and multiple treatment thresholds (as opposed to a single threshold), I employ a slightly modified version of the RD designs used by Urquiola and Verhoogen (2009) and Lacetera et al. (2012). I implement the RD design with the following regression specification:

$$y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f\left(Price_{it}\right) + \varepsilon_{it} \tag{4}$$

where y_{it} is some outcome for stock i on date t; $Ticksize_{it}$ is the discrete tick size; and $f\left(Price_{it}\right)$ is a flexible function of the stock price. If specified correctly, $f\left(Price_{it}\right)$ will capture all dependence of y_{it} and $Ticksize_{it}$ on the stock price away from the tick size price thresholds, such that the coefficient τ is estimated using only the variation in the tick size that occurs at the exact stock price levels where the tick size changes (the tick size discontinuities in panel a, Figure 2). The coefficient τ can be interpreted as the causal effect of tick sizes on y_{it}, under the identifying assumption that stocks are comparable in both their observable and unobservable stock characteristics at the price thresholds.

Consistent estimation of τ requires an assumption about the functional form of the relationship between y_{it} and the stock price. The RD literature has proposed two main approaches to estimating equation 4 when the functional form of this relationship is unknown. The first approach is to restrict the sample size on either side of a treatment threshold and estimate non-parametric local linear regressions around the threshold. The second approach, in contrast, involves using all the available data and selecting a flexible parametric specification for $f\left(Price_{it}\right)$.

While the local linear regression approach is theoretically more appealing (Hahn et al. 2001, Lee and Lemieux 2010), I follow Lacetera et al. (2012) and Urquiola and Verhoogen (2009) and estimate the regression discontinuity design globally by allowing for a flexible parametric specification of $f\left(Price_{it}\right)$. Following Lacetera et al. (2012), I approximate $f\left(Price_{it}\right)$ with a seventh order polynomial. The reason why I choose the parametric approach instead of the non-parametric local linear approach, is that my empirical setting departs from the 'standard' RD setting since there are multiple price thresholds that determine tick sizes. Instead of treating each tick size price threshold individually with local linear regressions, for convenience, I estimate the combined impact of all the thresholds within the same regression specification. The parametric approach yields the added benefit of allowing me to utilize more of the data which may improve statistical precision.

Stock prices may be more likely to cross a tick size price threshold on days when prices are volatile. To control for the influence of market-wide movements that can induce tick size changes, I add to equation 4 a full set of time fixed effects (α_t). Moreover, to control for unobserved and unchanging characteristics of a given stock, I add a full set of stock fixed effects to equation 4 (α_i). As a consequence, the identifying variation that is captured by the τ coefficient arises from stocks that cross a tick size price threshold at least once during the sample period, either from above or below.[9]

In the appendix of this article, I expose the regression discontinuity design to several validity tests and robustness specifications. The appendix shows that the main results are fairly stable across alternative polynomial specifications of $f\left(Price_{it}\right)$, and that the main results are robust to the inclusion of control variables. Finally, the appendix tests for and rejects discontinuities in y_{it} at placebo tick size price thresholds (price levels that do not affect the tick size).

In all regression specifications, standard errors are clustered at the stock-level.

[9]This is similar in spirit to the much-used difference-in-differences identification approach. The difference-in-differences estimator is measured as the change in outcomes for a treated group of stocks before and after an event relative to the corresponding change in outcomes for a control group of stocks unaffected by the event. Unlike the difference-in-differences approach, however, the regression discontinuity design in equation 4 only uses variation in outcomes that is generated on the exact dates when the tick size changes.

4.1 Summary of price threshold crossings

The identifying variation in equation 4 arises from stocks that cross tick size price thresholds either from above or below. Table 4 summarizes the occurrence of crossings of the NOK10, NOK15, NOK50, and NOK100 tick size price thresholds throughout the sample period 2008–2010. The table reports threshold crossings separately for non-OBX and OBX index stocks. For non-OBX stocks, there are 2330 tick size threshold crossings distributed across 157 unique stocks. The most-crossed price thresholds are NOK10 and NOK15, totalling more than 800 crossings (from above and below) for each threshold. The least-crossed price threshold is, by far, NOK100 with less than 200 crossings throughout the sample period.

For the OBX sample, there are 345 crossings of the 10NOK, 15NOK, 50NOK, and 100NOK price thresholds distributed across 26 unique stocks. Notice, however, that the actual number of tick size changes for OBX index stocks is less than the 345 price threshold crossings reported in Table 4. Due to a change in the tick size schedule for OBX index stocks in September 2009, crossings of the 10NOK (15NOK) price threshold in the first (second) half of the sample period 2008 – 2011 did not lead to tick size changes. In the empirical analysis, I account for the change in tick size schedules by estimating the regression discontinuity design separately for observations before and after September 2009.

5 Main results

In this section, I use a regression discontinuity design to estimate the causal impact tick size changes on the stock liquidity and trading volume at the Oslo Stock Exchange. The section begins by exploring the impact of tick size changes for liquid stocks in the OBX index, before it describes how the impact of tick sizes depends on initial stock liquidity.

5.1 Tick sizes in liquid stocks

The empirical results in Section 3 suggested that reducing the tick size for the most liquid stocks at the OSE (OBX index stocks) results in narrower bid-ask spreads, lower order book depth, and reduced trading volume. The conclusions in Section 3, however, arise from a before-and-after event study surrounding a single tick size reduction. Table 6, instead, uses the regression discontinuity design described in Section 4 to evaluate the causal impact of tick sizes on stock liquidity and trading volume for liquid stocks. The table presents estimates from the regression discontinuity design applied separately to two time periods: January 2008 – May 2009, and September 2009 – December 2011.[10]

Table 6 confirms that increasing the tick size for liquid stocks results in wider spreads and deeper order books. In the latest time period, September 2009 – December 2011, there is also weak evidence that increasing the tick size causes more trading volume. This effect, however, is not present in the earliest time period (January 2008 – May 2009), which suggests that the tick size, over time, may have become a more important factor for large-cap stock trading volume. A potential explanation for the increasingly benign impact of tick sizes on trading volume could be the recent explosion in high-frequency trading (HFT), both at the OSE (Jørgensen et al. 2016) and around the world in general. Recent empirical work by O'Hara et al. (2015) suggests that HFTs prefer to trade in large-tick size environments, since large tick sizes exacerbate the HFT speed advantage. The interaction between an increase in HFT activity and their presumed preference for large-tick trading may explain why larger tick sizes improve trading volume in the latest time period (September 2009 – December 2011) but not the earliest time period (January 2008 – May 2009).

The results in Table 6 not only validate the before-and-after estimates from Section 3; they also line up with the existing empirical tick size literature. A voluminous literature, predominantly focusing on regulatory tick size changes using before-and-after estimators, has established that increasing the tick size leads to wider bid-ask spreads and deeper order books (see for example the recent survey by the Securities and Exchange Commission 2012). My results complement the existing literature by showing that the established relationships between tick sizes, bid-ask spreads, and order book depths for liquid stocks are robust to a rigorous regression discontinuity design. Moreover, my results add to the existing empirical literature by showing a potentially time-varying relationship between tick sizes and trading volume for liquid stocks.

[10]The overall sample period ($2008 - 2011$) is split into two separate periods to account for the change in the tick size schedule for OBX index stocks in late August 2009. Table 1 provides detailed information on the tick size schedules used in the periods January 2008 – May 2009, and September 2009 – December 2011.

5.2 Tick sizes in illiquid stocks

Section 5.1 established a strong effect of increasing the tick size on stock liquidity for liquid stocks at the Oslo Stock Exchange. Motivated both by recent theoretical predictions by Buti et al. (2015) and by the current tick size policy debate in the United States, I turn to explore whether tick sizes affect the market quality of liquid and illiquid stocks differently. In order to estimate such cross-sectional treatment effect heterogeneity, I employ the sample of non-OBX index stocks in the period $2008-2010$. Unlike the OBX sample, which only holds at most 25 stocks, the non-OBX sample comprises a large number of both liquid and illiquid stocks, which is a prerequisite for exploring cross-sectional heterogeneity.

I begin by assessing the average impact of increasing the tick size for the full sample of non-OBX index stocks using a regression discontinuity design and data from the time period $2008-2011$. The bottom panel of Table 6 shows that the average effect of increasing the tick size for the full sample of non-OBX index stocks is to widen spread measures of liquidity and to improve order book depth. At the same time, I find no relationship between trading volume and tick sizes for this sample of stocks. Thus, the average effect of increasing the tick size for non-OBX stocks does not appear to differ much from the average effect of increasing the tick size for liquid stocks (top two panels of Table 6).

The average treatment effects displayed in Table 6, however, may conceal considerable heterogeneity. To further explore how the effect of tick sizes depends on initial stock liquidity, I split the sample of non-OBX stocks into equally-sized terciles based on stock trading volume. For each stock in the non-OBX sample, I compute the average trading volume in January 2008 (the first month in the sample period). Stocks are then sorted into terciles based on the January 2008 trading volume. The tercile a stock belongs to remains the same throughout the sample period $2008-2011$. Moreover, in the upcoming empirical analyses I only use data from February 2008 and onwards. This procedure ensures that the tercile formation itself cannot be affected by tick size changes.

Table 5 presents descriptive statistics which illustrate the variation in stock characteristics that the trading volume terciles capture. First, although the terciles are formed based on a single liquidity metric — stock trading volume — Table 5 shows that the tercile formation could equally well have been based on any other market quality metric. Specifically, the table shows that Tercile 1 (least traded) consistently has the widest spreads, most shallow books, highest volatility, and (naturally) the lowest trading volumes. For example, the median trading volume in Tercile 1 is 80000NOK (1 USD $\approx$ 8 NOK), the median order book depth is 75 000NOK, and the median relative quoted spread is almost 5% of the current midquote. Tercile 3 (most traded), in contrast, represents a reasonably liquid trading environment, with a median trading volume of almost two million NOK, a median order book depth of 162 000NOK, and a median relative quoted spread of 1.3% of the current midquote. Indeed, along some dimensions of stock liquidity, such as order book depth and transaction costs, trading in Tercile 3 stocks appears comparable to the statistics of the liquid large-cap stocks in Table 2.

Second, the trading volume terciles capture variation in how constrained the bid-ask spread is by the tick size — a variation that is potentially important for understanding the empirical results. O'Hara et al. (2015) explore whether changes to the relative tick size affect stocks in a one-tick environment (the bid-ask spread is equal to the tick size) and stocks in multi-tick environments differently. They show that in the one-tick environment, an increase in the relative tick size leads to more trading volume and more order book depth. In contrast, in the multi-tick environment an increase in the relative tick size leads to lower trading volume and less order book depth. Table 5 shows that stocks in Tercile 3 tend to trade close to a one-tick environment, with a median ticks-per-spread of only 3. In contrast, Terciles 2 and 3 tend to trade in a multi-tick environment, with median ticks-per-spread of 6 and 10 respectively.

Finally, the trading volume terciles capture a variation in market capitalization. The average market capitalization is monotonically increasing in the terciles, from 740 million NOK in Tercile 1 to 1505 million NOK in Tercile 2 and finally 2716 million NOK in Tercile 3. For comparison, the eligibility criteria for the recently implemented Tick Size Pilot Program in the United States is that stocks should have a market capitalization of less than \$3 (approximately 18 billion NOK). Clearly, judged by this criteria alone, the average stock in all the trading volume terciles would be eligibile for the tick size pilot.

Table 7 presents estimates from the regression discontinuity design applied separately to each of the trading volume terciles. For the most illiquid stocks (Tercile 1), there are no measurable effects of increasing the tick size on the quality of trading. Specifically, increasing the tick size for this sample of stocks does not affect spread measures of liquidity, order book depth, or trading volume. In contrast, for Tercile 3 (most traded), increasing the tick size causes significantly wider spreads and deeper order books, suggesting that the average effect for non-OBX stocks in Table 6 is primarily driven by the most liquid stocks in the distribution.

Splitting the sample into terciles provides a somewhat coarse insight into how the effect of tick sizes differs depending on initial stock liquidity. As an alternative approach to illustrate treatment effect heterogeneity, I split the sample into quantiles instead terciles, using the same ranking procedure as before. Table 8 presents estimates of the regression discontinuity design applied separately to each of the quantiles. The table confirms the impression from Table 7. For the bottom 60% of the liquidity distribution, I find no effects of increasing the tick size on either liquidity or trading volume. Instead, for the top 40% of the liquidity distribution there is a strong and statistically significant impact of tick size changes on both spreads and order book depths, but no impact on volatility or trading volume.

6 Discussion and concluding remarks

Estimates from a so-called regression discontinuity design reveal that the causal effect of increasing the tick size, the minimum price increment on a stock exchange, differs depending on the initial stock liquidity. For liquid stocks at the Oslo Stock Exchange in the period $2008 - 2011$, increasing the tick size leads to wider bid-ask spreads and deeper order books, and has a weakly significant and potentially time-varying positive impact on trading volume. For the most illiquid stocks at the Oslo Stock Exchange, however, changing the tick size has no impact on bid-ask spreads, order book depths, volatility, or trading volume.

There are several implications of the results in this paper. First, my empirical results have implications for the current theoretical debate over the potentially heterogeneous impact of tick sizes on stocks with different liquidity. Buti et al. (2015) predict that increasing the tick size for illiquid stock may improve stock liquidity and decrease trading volume. My results provide little empirical support for this prediction. Meanwhile, my results suggest that increasing the tick size for liquid stocks may in fact increase both order book depth and widen the bid-ask spread, while at the same time increasing trading volume. These results are largely consistent with the theoretical predictions by Buti et al. (2015) for liquid order books.

Second, the recently implemented "Tick Size Pilot Program" in the United States, which has increased the tick size for a large number of small and medium sized firms, reflects a similar suspicion that the "one size fits all" penny tick size in the United States may not be optimal for the entire distribution of firms. The main argument behind the tick size pilot is that small tick sizes may be optimal for liquid (large-cap) securities, as it will reduce trading costs, while large tick sizes may be optimal for illiquid (small-cap) securities, as it will provide incentives for liquidity provision in these stocks and therefore enhance overall trading volume. The results in this paper suggest that smaller tick sizes may reduce transaction costs for liquid stocks, however only at the expense of reduced order book depth. For illiquid stocks, however, such a trade-off does not exist as the tick size does not appear to affect any measure of small-cap market quality. Thus, my estimates suggest that other market structure tools than tick sizes are needed if the object is to improve the quality of trading in illiquid stocks.

Third, the results in this paper illustrate the importance of evaluating heterogeneous responses to equity market policy changes. I show that tick size changes appear to have heterogeneous effects across the stock liquidity distribution — a large portion of the liquidity distribution experiences no effect from tick size changes (illiquid stocks) while a small portion of the liquidity distribution experiences a considerable effect from tick size changes (liquid stocks). Nevertheless, the resulting average treatment effect, which is estimated across the entire distribution of stocks, is measured to be highly statistically significant. In terms of policy advice and extrapolation to alternative contexts, this average treatment effect may be seriously misleading when not accompanied with information about the underlying effect heterogeneity.

Meanwhile, I also caution about the interpretation of the results in the present paper. Illiquid stocks are, in my setting, defined jointly by their low trading volume, shallow order books, high transaction costs, and their unconstrained bid-ask spreads. This joint definition of illiquidity is not by purposeful design, but is rather an artifact of significant correlation between liquidity measures — differentiating stocks on one liquidity measure typically implies differentiating on another liquidity measure as well. For this reason, I cannot determine whether heterogeneity in the effect of tick sizes is driven primarily by any specific liquidity measure, or simply by the combination of all the liquidity measures. In their theoretical model, Buti et al. (2015) define illiquid stocks exclusively based on order book depth. My empirical analysis cannot, therefore, be interpreted as a direct test of the theoretical predictions in Buti et al. (2015). Instead, my empirical results can be interpreted as showing that the effect of tick sizes varies depending on a more general definition of stock liquidity.

References

BACIDORE, J., R. H. BATTALIO, AND R. H. JENNINGS (2003): "Order submission strategies, liquidity supply, and trading in pennies on the New York Stock Exchange," *Journal of Financial Markets*, 6, 337 – 362.

BUTI, S., F. CONSONNI., B. RINDI., Y. WEN, AND I. M. WERNER (2015): "Tick Size: Theory and Evidence," Working Paper. Rotman School of Management Working Paper No. 2485069; Fisher College of Business Working Paper 2015-03-04; Charles A. Dice Center Working Paper No. 2015-04. Available at SSRN: http://ssrn.com/abstract=2485069.

CHAKRAVARTY, S., R. A. WOOD, AND R. A. VAN NESS (2004): "Decimals And Liquidity: A Study Of The Nyse," *Journal of Financial Research*, 27, 75–94.

GOLDSTEIN, M. A. AND K. A. KAVAJECZ (2000): "Eights, Sixteenths and Market Depth: Changes in tick size," *Journal of Financial Economics*, 56, 125–149.

HAHN, J., P. TODD, AND W. VAN DER KLAAUW (2001): "Identification and estimation of treatment effects with a regression-discontinuity design," *Econometrica*, 69, 201–209.

JØRGENSEN, K., J. SKJELTORP, AND B. A. ØDEGAARD (2016): "Throttling hyperactive robots - Order to Trade Ratios at the Oslo Stock Exchange," Working Paper, University of Stavanger.

LACETERA, N., D. G. POPE, AND J. R. SYDNOR (2012): "Heuristic Thinking and Limited Attention in the Car Market," *American Economic Review*, 102, 2206–36.

LEE, D. S. AND T. LEMIEUX (2010): "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 48, 281–355.

MCCRARY, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 142, 698 – 714, the regression discontinuity design: Theory and applications.

MELING, T. G. (2016): "Anonymous trading in equities," Working Paper. Available at SSRN: https://ssrn.com/abstract=2656161.

MELING, T. G. AND B. A. ØDEGAARD (2016): "Tick Size Wars," Working Paper. Available at SSRN: https://ssrn.com/abstract=2866943.

O'HARA, M., G. SAAR, AND Z. ZHO (2015): "Relative Tick Size and the Trading Environment," Working Paper, SSRN.

RONEN, T. AND D. G. WEAVER (2001): "Teenes' Anyone?" *Journal of Financial Markets*, 4, 231–260.

SECURITIES AND EXCHANGE COMMISSION (2012): "Report to Congress on Decimalization," Available at http://www.sec.gov/news/studies/2012/decimalization-072012.pdf.

URQUIOLA, M. AND E. VERHOOGEN (2009): "Class-Size Caps, Sorting, and the Regression-Discontinuity Design," *American Economic Review*, 99, 179–215.

7 Figures

Figure 1: Stock price density

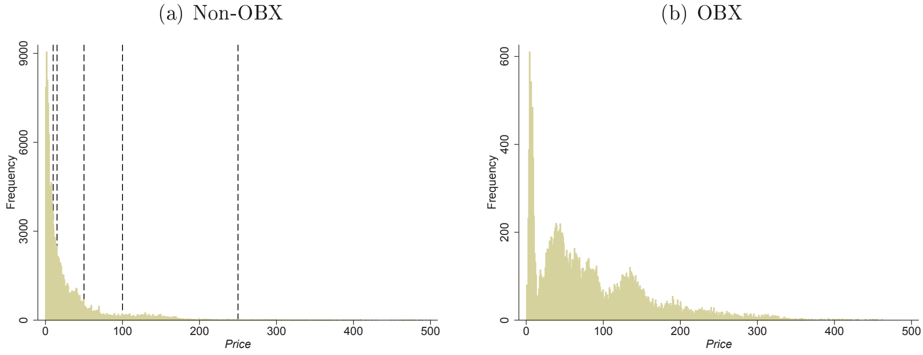

Note: The figure presents a histogram of stock prices in the period 2008 − 2011, separately for non-OBX index stocks (panel a) and OBX index stocks (panel b). Each bar has a width of 1NOK, and there are 500 bars. Vertical breaks indicate price levels where the tick size increases. Vertical breaks have been excluded from panel (b) because the tick size schedule for OBX index stocks changed during the time period 2008 − 2011.

Figure 2: Illustration of regression discontinuity design

(a) Tick size over price

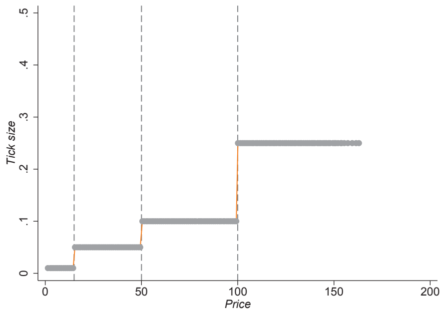

(b) Log(Effective spread) over price

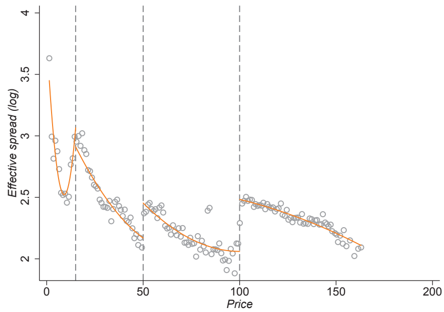

Note: The figure illustrates the regression discontinuity design. Panel (a) plots tick sizes as a function of prices. Panel (b) plots effective spreads (log) as a function of prices. Both panels (a) and (b) plot observations from a sample of OBX index stocks in the period January 2008 – May 2009. Vertical dashed lines indicate stock price levels where the tick size increases. The regression discontinuity design estimates the impact of tick sizes as the discontinuous change in outcomes at the exact price levels where the tick size changes.

8 Tables

Table 1: Tick size schedules

OBX index stocks					
June, 2003 - August, 2006		*Sept. 2006 - May, 2009*		*Sept. 2009 - Dec. 2012*	
Price band	**Tick size**	**Price band**	**Tick size**	**Price band**	**Tick size**
0 - 9.99	0.01	0 - 14.99	0.01	0 - 0.4999	0.0001
10 - 49.9	0.10	15 - 49.95	0.05	0.5 - 0.9995	0.0005
50 - 149.75	0.25	50 - 99.9	0.1	1 - 4.9990	0.001
150 - 999.5	0.50	100 - 249.75	0.25	5 - 9.995	0.005
1000 -	1.00	250 - 499.50	0.5	10 - 49.990	0.01
		>500	1	50 - 99.95	0.05
				100 - 499.90	0.1
				500 - 999.50	0.5
				1000 - 4999	1
				5000 - 9995	5
				>10000	10
Non-OBX index stocks					
June, 2003 - August, 2006		*Sept. 2006 - May, 2009*		*Sept. 2009 - Dec. 2012*	
Price band	**Tick size**	**Price band**	**Tick size**	**Price band**	**Tick size**
0 - 9.99	0.01	0 - 9.99	0.01	0 - 9.99	0.01
10 - 49.9	0.10	10 - 14.95	0.05	10 - 14.95	0.05
50 - 149.75	0.25	15 - 49.9	0.1	15 - 49.9	0.1
150 - 999.50	0.50	50 - 99.75	0.25	50 - 99.75	0.25
1000 -	1.00	100 - 249.5	0.5	100 - 249.5	0.5
		>250	1	>250	1

Note: The table shows the evolution of tick size schedules at OSE. Tick sizes are determined by stock price bands (in NOK). The top panel shows tick size schedules for large-cap stocks. The bottom panel shows tick size schedules for small-cap stocks.

Table 2: Summary statistics

	μ	σ	Min.	Median	Max.	N
OBX stocks						
Market cap. (mNOK)	29200.59	38028.17	0.00	13426.00	217595.56	16555
Relative spread (bps)	29.47	20.54	4.09	24.45	461.52	16554
Realized spread (bps)	2.21	6.98	-53.08	1.46	640.16	16548
Depth (thousands NOK)	737.77	1451.47	18.08	433.75	45649.04	16554
Realized volatility (pp)	0.74	1.53	0.02	0.45	46.86	16542
Volume (mNOK)	155.71	223.88	0.01	97.73	13443.04	16555
Stock price	51.56	38.07	0.38	45.04	165.00	16555
Tick size	0.05	0.06	0.00	0.05	0.25	16555
Relative tick size	0.10	0.09	0.02	0.08	1.03	16555
Ticks-per-spread	3.69	5.59	0.00	2.00	207.50	16553
non-OBX stocks						
Market cap. (mNOK)	1799.81	3080.58	0.00	857.24	105517.95	124770
Relative spread (bps)	369.42	295.47	11.49	286.71	1497.12	136800
Realized spread (bps)	59.74	137.24	-1740.74	26.68	2352.94	93227
Depth (thousands NOK)	348.30	4425.28	0.06	105.72	481635.51	145325
Realized volatility (pp)	1.03	1.52	0.02	0.70	49.30	75209
Volume (mNOK)	4.77	65.80	0.00	0.43	17368.01	124759
Stock price	23.87	30.90	0.02	12.00	198.50	124770
Tick size	0.08	0.11	0.01	0.05	0.50	124770
Relative tick size	0.57	1.76	0.10	0.36	50.00	124770
Ticks-per-spread	10.44	18.13	0.00	5.00	1600.00	123924

Note: The table gives summary statistics for OBX index stocks (large-caps) and non-OBX index stocks (small-caps) at the Oslo Stock Exchange (OSE) in the period 2008 – 2011. The stock characteristics are market capitalization (millions NOK); relative and realized spreads (basis points); order book depth (thousands NOK); realized volatility (millions NOK); stock price (NOK); tick size (NOK); relative tick size (tick size relative to stock price); ticks-per-spread (tick size relative to quoted spread). The table lists means (μ), standard deviations (σ), minimum (Min.) and maximum values (Max.), medians, and number of observations (N).

Table 3: Before-and-after estimates of tick size reduction

	Dependent variable				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
β	-0.10*	-0.07	-0.42***	0.00	-0.12*
	(-2.27)	(-0.53)	(-9.36)	(0.82)	(-2.14)
N	500	332	500	500	500
Adj. R^2	0.01	-0.00	0.07	-0.00	0.00

Note: The table gives before-and-after estimates of the impact of the July 6, 2009 tick size reduction for OBX index stocks at the Oslo Stock Exchange. The regression specification is $y_{it} = \alpha + \beta Post_t + \epsilon_{it}$, where $Post_t = 1$ for observations after July 6, 2009. The sample comprises ten trading days before and ten trading days after July 6, 2009 for all the 25 stocks in the OBX index. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. Standard errors clustered at the stock-level. t-statistics in parentheses.
* $p < 0.05$, ** $p < 0.025$, *** $p < 0.01$

Table 4: Tick size price threshold crossings

	Non-OBX sample in 2008-2011 Number of price threshold crossings
NOK10 from below	408
NOK15 from below	387
NOK50 from below	236
NOK100 from below	92
NOK10 from above	437
NOK15 from above	428
NOK50 from above	247
NOK100 from above	95
Unique stocks crossing any threshold	157
	OBX sample in 2008-2011 Number of price threshold crossings
NOK10 from below	28
NOK15 from below	13
NOK50 from below	68
NOK100 from below	54
NOK10 from above	34
NOK15 from above	17
NOK50 from above	74
NOK100 from above	57
Unique stocks crossing any threshold	26

Note: The table summarizes the occurrence of tick size price threshold crossings in the sample period 2008–2011 for non-OBX index stocks. Threshold crossings are summarized separately for crossings from below and above a price threshold. Threshold crossings are defined at the daily level using end-of-day prices.

Table 5: Summary statistics: Liquidity terciles

	μ	σ	Min.	Median	Max.	N
Tercile 1 (Least traded)						
Market cap. (mNOK)	740.67	827.30	2.10	498.37	6075.00	27069
Relative spread (bps)	556.27	321.79	26.49	483.03	1497.12	34092
Realized spread (bps)	85.22	158.30	-1740.74	48.10	1698.83	13236
Depth (thousands NOK)	151.24	772.95	0.59	75.61	40552.00	38950
Realized volatility (pp)	1.19	1.40	0.02	0.83	36.02	7868
Volume (mNOK)	1.14	15.05	0.00	0.08	1333.67	27009
Stock price	17.18	18.46	0.08	10.85	129.00	27069
Tick size	1.20	1.31	0.20	1.00	10.00	27069
Relative tick size	0.51	0.70	0.10	0.38	12.50	27069
Ticks-per-spread	16.68	22.57	0.40	10.00	700.00	26695
Tercile 2						
Market cap. (mNOK)	1505.67	2426.09	0.95	642.97	65955.11	40655
Relative spread (bps)	384.22	259.14	12.07	323.17	1497.12	44120
Realized spread (bps)	76.25	159.88	-1172.24	39.05	2248.06	29326
Depth (thousands NOK)	213.27	744.75	1.85	95.13	35173.72	45449
Realized volatility (pp)	1.11	1.44	0.02	0.75	35.56	21100
Volume (mNOK)	2.41	18.07	0.00	0.25	1338.78	40814
Stock price	27.15	36.44	0.02	12.30	198.50	40655
Tick size	1.98	2.70	0.20	1.00	10.00	40655
Relative tick size	0.61	1.65	0.10	0.37	50.00	40655
Ticks-per-spread	10.62	15.22	0.20	6.00	400.00	40478
Tercile 3 (Most traded)						
Market cap. (mNOK)	2716.11	4094.54	6.64	1798.84	105517.95	46955
Relative spread (bps)	205.87	190.54	11.49	137.13	1482.50	47309
Realized spread (bps)	41.97	111.66	-847.46	19.01	1666.67	44332
Depth (thousands NOK)	632.43	7553.39	4.50	162.43	481635.50	47809
Realized volatility (pp)	0.98	1.62	0.02	0.66	49.30	41335
Volume (mNOK)	9.11	55.58	0.00	2.03	6034.40	46858
Stock price	25.04	32.65	0.02	11.00	198.50	46955
Tick size	1.75	2.39	0.00	1.00	10.00	46955
Relative tick size	0.63	2.36	0.02	0.35	50.00	46955
Ticks-per-spread	6.03	12.11	0.00	3.00	1070.00	46882

Note: The table gives summary statistics from the time period $2008-2011$ separately for terciles that are formed based on average trading volume in January 2008. The stock characteristics are market capitalization (millions NOK); relative and realized spreads (basis points); order book depth (thousands NOK); realized volatility (millions NOK); stock price (NOK); tick size (NOK); relative tick size (tick size relative to stock price); ticks-per-spread (tick size relative to quoted spread). The table lists means (μ), standard deviations (σ), minimum (Min.) and maximum values (Max.), medians, and number of observations (N).

Table 6: Main results

	OBX sample in period: 2008 - May 2009				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.20***	0.39***	0.20***	0.16*	-0.01
	(7.43)	(5.18)	(3.42)	(2.29)	(-0.16)
N	6530	4986	6530	6518	6530
Adj. R^2	0.81	0.21	0.82	0.12	0.71
	OBX sample in period: September 2009 - 2011				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.12	0.58***	0.73***	0.10*	0.22**
	(1.08)	(3.80)	(5.55)	(2.29)	(2.67)
N	10023	7411	10023	10024	10025
Adj. R^2	0.87	0.33	0.90	0.11	0.82
	non-OBX sample in period: 2008 - 2011				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.07***	0.13***	0.10***	0.04*	0.01
	(4.54)	(6.11)	(4.16)	(2.17)	(0.17)
N	121296	82093	123954	74997	124770
Adj. R^2	0.66	0.37	0.51	0.17	0.51

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for stocks at the Oslo Stock Exchange in the period 2008 – 2011. The regression discontinuity design is run separately for OBX index stocks and non-OBX index stocks. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f\,(price_{it}) + \varepsilon_{it}$, where $f\,(price_{it})$ is a 7th order polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. The coefficient τ identifies discrete jumps in y_{it} at the exact stock price levels where the tick size changes. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The τ regression coefficient has been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses.

* p < 0.05, ** p < 0.025, *** p < 0.01

Table 7: Tercile regressions

	Tercile regressions: 2008-2011				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
Tercile 1 (Least traded)					
τ	0.05	-0.02	0.06	0.10	0.06
	(1.86)	(-0.45)	(1.35)	(1.59)	(0.66)
N	24762	10856	25905	7737	26344
Adj. R^2	0.42	0.26	0.38	0.24	0.26
Tercile 2					
τ	0.05	0.07	0.01	-0.01	-0.06
	(1.80)	(1.97)	(0.42)	(-0.25)	(-0.93)
N	38657	24602	39364	20206	39484
Adj. R^2	0.55	0.29	0.42	0.25	0.33
Tercile 3 (Most traded)					
τ	0.09***	0.18***	0.17***	0.06*	0.05
	(4.57)	(5.40)	(3.90)	(2.13)	(1.42)
N	43742	37552	44080	38495	44101
Adj. R^2	0.61	0.32	0.53	0.15	0.47

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for stocks at the Oslo Stock Exchange in the period February 2008 to December 2011. The regression discontinuity design is run separately for terciles that are formed based on average trading volume in January 2008. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f\,(price_{it}) + \varepsilon_{it}$, where $f\,(price_{it})$ is a 7th order polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. The coefficient τ identifies discrete jumps in y_{it} at the exact stock price levels where the tick size changes. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The τ regression coefficient has been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses.

** p < 0.05, ** p < 0.025, *** p < 0.01*

Table 8: Quantile regressions

	Quantile regressions: 2008-2011				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
Quantile 1 (Least traded)					
τ	0.04	-0.06	0.03	0.26**	-0.05
	(1.56)	(-0.70)	(0.45)	(2.46)	(-0.41)
N	11978	4674	12854	3006	13244
Adj. R^2	0.42	0.26	0.41	0.26	0.23
Quantile 2					
τ	-0.02	-0.05	0.07	0.02	0.15
	(-0.43)	(-1.05)	(1.67)	(0.36)	(1.62)
N	20839	10905	21375	8299	21439
Adj. R^2	0.44	0.26	0.35	0.24	0.26
Quantile 3					
τ	0.04	0.03	-0.00	-0.06	-0.05
	(1.22)	(0.50)	(-0.05)	(-0.73)	(-0.63)
N	24152	14962	24462	12143	24556
Adj. R^2	0.52	0.25	0.43	0.22	0.30
Quantile 4					
τ	0.09***	0.12***	0.14***	0.05	0.05
	(3.17)	(3.47)	(3.02)	(1.64)	(0.65)
N	25528	20170	25781	19300	25808
Adj. R^2	0.58	0.27	0.46	0.21	0.38
Quantile 5 (Most traded)					
τ	0.08***	0.18***	0.16***	0.04	0.05
	(3.00)	(4.73)	(2.96)	(1.39)	(1.54)
N	24664	22258	24877	23577	24882
Adj. R^2	0.60	0.36	0.53	0.15	0.48

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for stocks at the Oslo Stock Exchange in the period February 2008 to December 2011. The regression discontinuity design is run separately for quantiles that are formed based on average trading volume in January 2008. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f\left(price_{it}\right) + \varepsilon_{it}$, where $f\left(price_{it}\right)$ is a 7th order polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. The coefficient τ identifies discrete jumps in y_{it} at the exact stock price levels where the tick size changes. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The τ regression coefficient has been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses.
* $p < 0.05$, ** $p < 0.025$, *** $p < 0.01$

A Robustness of regression discontinuity design

This section explores the sensitivity of the regression discontinuity design to alternative specifications. As described in Section 4, I implement the RD design with the following regression specification:

$$y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f\left(Price_{it}\right) + \varepsilon_{it} \tag{5}$$

where y_{it} is some outcome for stock i on date t; $Ticksize_{it}$ is the discrete tick size; and $f\left(Price_{it}\right)$ is a flexible function of the stock price. If specified correctly, $f\left(Price_{it}\right)$ captures all dependence of y_{it} and $Ticksize_{it}$ on the stock price away from the tick size price thresholds, such that the coefficient τ is estimated using only the variation in the tick size that occurs at the exact stock price levels where the tick size changes. The coefficient τ can be interpreted as the causal effect of tick sizes on y_{it}, under the identifying assumption that stocks are comparable in both their observable and unobservable stock characteristics at the price thresholds.

This section modifies equation 5 threefold. First, Section A.1 allows for a variety of different polynomial specifications of $f\left(Price_{it}\right)$. Second, Section A.2 tests for discontinuities in y_{it} at placebo tick size price thresholds. Third, Section A.3 adds control variables to equation 5. All the robustness tests are based on the sample of non-OBX index stocks in the period $2008 - 2011$. Therefore, the results in this section can be compared to the baseline results presented in the bottom panel of Table 6.

A.1 Alternative polynomial specifications

The regression discontinuity specification in Section 5 assumes that the relationship between stock prices and outcomes can be adequately captured by a seventh order polynomial. In Table 9, however, I relax this assumption and explore the robustness of the RD design to alternative polynomial specifications. The table estimates equation 5 separately for linear, quadratic, cubic, quartic, and quintic specifications of $f\left(\right)$. Table 9 shows that the estimates of τ remain fairly stable across polynomial specifications.

Table 9: Alternative polynomial specifications

	Polynomial specification				
	Linear	Quadratic	Cubic	Quartic	Quintic
Relative spread	0.04***	0.06***	0.05***	0.06***	0.07***
	(3.12)	(4.44)	(3.90)	(4.54)	(4.92)
N	121296	121296	121296	121296	121296
Adj. R^2	0.65	0.65	0.65	0.65	0.65
Realized spread	0.10***	0.11***	0.10***	0.13***	0.15***
	(5.19)	(6.83)	(5.64)	(6.46)	(7.01)
N	82093	82093	82093	82093	82093
Adj. R^2	0.36	0.36	0.36	0.37	0.37
Depth	0.11***	0.08***	0.11***	0.10***	0.10***
	(4.19)	(3.09)	(4.13)	(4.17)	(4.09)
N	123954	123954	123954	123954	123954
Adj. R^2	0.49	0.51	0.51	0.51	0.51
Volatility	0.02	0.04***	0.02	0.04**	0.06***
	(1.46)	(2.74)	(1.40)	(2.44)	(2.88)
N	74997	74997	74997	74997	74997
Adj. R^2	0.16	0.17	0.17	0.17	0.17
Trading volume	0.02	0.00	0.01	0.01	0.00
	(0.64)	(0.03)	(0.38)	(0.33)	(0.05)
N	124770	124770	124770	124770	124770
Adj. R^2	0.51	0.51	0.51	0.51	0.51

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for non-OBX index stocks at the Oslo Stock Exchange in the period 2008 – 2011. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f(price_{it}) + \varepsilon_{it}$, where $f(price_{it})$ represents a flexible polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. The regression specification is estimated separately for linear through cubic polynomial specifications of $f(price_{it})$. The coefficient τ identifies discrete jumps in y_{it} at the exact stock price levels where the tick size changes. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The top panel includes the natural logarithm of market capitalization as a control variable. The bottom panel includes both market capitalization and $y_{i,t-1}$ as control variables. The τ regression coefficient has been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses. ** p < 0.05, ** p < 0.025, *** p < 0.01*

A.2 Placebo thresholds

The identifying assumption in the regression discontinuity design is that the stock price thresholds that increase or decrease the tick size only affect stock outcomes through their impact on the tick size. To assess whether or not this assumption is plausible, I explore whether the effects documented in Section 5 are exclusive to stock price thresholds that mandate tick size changes. To this end, I generate a 'placebo' tick size variable which starts at 0.01NOK and increases to 0.1NOK, 0.2NOK, 0.3NOK, and 0.4NOK at the 25NOK, 50NOK, 75NOK, and 125NOK price thresholds, respectively.

In Table 10, I estimate the regression discontinuity design with both the actual tick size variable ($Ticksize_{it}$) and the 'placebo' tick size variable ($Ticksize_{it}^{Placebo}$) as explanatory variables. Table 10 shows that the estimated impact of $Ticksize_{it}$ remains similar to estimates from the baseline specification. Reassuringly, the table also shows that for all outcome variables except for realized volatility, there is no impact of $Ticksize_{it}^{Placebo}$.

Table 10: Placebo tests

	Dependent variable				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.06***	0.13***	0.09***	0.02	-0.00
	(4.66)	(5.88)	(3.87)	(1.62)	(-0.04)
$\tau^{Placebo}$	-0.01	-0.02	-0.05	-0.07*	-0.03
	(-0.34)	(-0.61)	(-1.48)	(-2.25)	(-0.52)
N	121296	82093	123954	74997	124770
Adj. R^2	0.66	0.37	0.51	0.17	0.51

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for non-OBX index stocks at the Oslo Stock Exchange in the period $2008-2011$. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + \tau^{Placebo} Ticksize_{it}^{Placebo} + f\left(price_{it}\right) + \varepsilon_{it}$, where $f\left(price_{it}\right)$ is a 7th order polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. $Ticksize_{it}$ is the tick size based on the actual tick size schedule and $Ticksize_{it}^{Placebo}$ is the tick size based on a fictional tick size schedule with tick size price thresholds at 25NOK, 75NOK, and 125NOK. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The τ and $\tau^{Placebo}$ regression coefficients have been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses.
** $p < 0.05$, ** $p < 0.025$, *** $p < 0.01$*

A.3 Control variables

If the regression discontinuity design is valid, there is no need to add control variables to equation 5 (Lee and Lemieux 2010). This is because randomness in whether a stock is priced marginally above or marginally below a tick size price threshold ensures that stocks on either side of the price threshold are comparable in their observable characteristics. Nevertheless, a common validity test in the regression discontinuity design literature is to estimate the RD design with non-outcome covariates as controls.

In my setting, there are few candidate control variables — most of the covariates can either be considered as outcomes (such as the liquidity measures) or the covariates do not vary on a sufficiently high frequency (such as earnings or assets). Nevertheless, there are two non-outcome daily frequency covariates that can be added to equation 5. The first is the natural logarithm of daily market capitalization. The second is the lagged outcome variable ($y_{i,t-1}$). As discussed in Lee and Lemieux (2010), adding the $y_{i,t-1}$ as a control may improve statistical precision when y_{it} is highly persistent.

In the top panel of Table 11, I estimate the regression discontinuity design using only market capitalization as a control variable. The estimates of τ from this specification are almost identical to the baseline specification. In the bottom panel of Table 11, I control for both market capitalization and $y_{i,t-1}$. Including $y_{i,t-1}$ in equation 5 reduces the magnitudes of the regression coefficients, but the statistical inference remains unchanged.

Table 11: Control variables

	Controlling for market cap.				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.06***	0.12***	0.11***	0.04**	0.03
	(4.49)	(6.28)	(4.59)	(2.56)	(1.02)
N	119223	80763	121744	73826	122471
Adj. R^2	0.69	0.39	0.54	0.18	0.53
	Controlling for market cap. and lagged outcome				
	Relspread	*Rspread*	*Depth*	*Volatility*	*Volume*
τ	0.02***	0.10***	0.04***	0.03**	0.03
	(4.88)	(6.02)	(5.07)	(2.28)	(1.23)
N	91792	50297	94843	49637	95297
Adj. R^2	0.83	0.43	0.77	0.29	0.56

Note: The table gives regression discontinuity estimates of the effect of increasing the tick size on market quality outcomes, for non-OBX index stocks at the Oslo Stock Exchange in the period 2008 − 2011. The regression specification is $y_{it} = \alpha_i + \alpha_t + \tau Ticksize_{it} + f(price_{it}) + \varepsilon_{it}$, where $f(price_{it})$ is a 7th order polynomial of the stock price and α_i and α_t are stock and time fixed effects, respectively. The coefficient τ identifies discrete jumps in y_{it} at the exact stock price levels where the tick size changes. *Relspread* is the relative bid-ask spread, log-transformed. *Rspread* is the realized spread, log-transformed. *Depth* is the order book depth, log-transformed. *Volatility* is realized volatility measured in percentages. *Volume* is the NOK trading volume, log-transformed. The top panel includes the natural logarithm of market capitalization as a control variable. The bottom panel includes both market capitalization and $y_{i,t-1}$ as control variables. The τ regression coefficient has been scaled, and can be interpreted as the change in y_{it} given a 0.05NOK increase in the tick size. Standard errors clustered at the stock-level. t-statistics in parentheses.

* $p < 0.05$, ** $p < 0.025$, *** $p < 0.01$