

This file was downloaded from BI Open Archive, the institutional repository (open access) at BI Norwegian Business School <http://brage.bibsys.no/bi>.

It contains the accepted and peer reviewed manuscript to the article cited below. It may contain minor differences from the journal's pdf version.

Jørgensen, K., Skjeltop, J., & Ødegaard, B. A. (2018). Throttling hyperactive robots: Order- to-trade ratios at the Oslo Stock Exchange. *Journal of Financial Markets*, 37(January), 1-16. DOI: <https://doi.org/10.1016/j.finmar.2017.09.001>

Copyright policy of Elsevier, the publisher of this journal.
The author retains the right to post the accepted author manuscript on open web sites operated by author or author's institution for scholarly purposes, with an embargo period of 0-36 months after first view online.
<http://www.elsevier.com/journal-authors/sharing-your-article#>



Throttling hyperactive robots – order-to-trade ratios at the Oslo Stock Exchange[☆]

Kjell Jørgensen¹, Johannes Skjeltop², Bernt Arne Ødegaard^{3,*}

Abstract

We investigate the effects of introducing a fee on excessive order-to-trade ratios (OTRs) on market quality at the Oslo Stock Exchange (OSE). We find that traders reacted to the regulation as measured OTRs fell. However, market quality, measured with depth, spreads, and realized volatility, remain largely unaffected. This result differs sharply from the experience in other markets, such as Italy and Canada, where similar regulatory changes have been accompanied by a worsening of liquidity. The unchanged market quality at the OSE is likely due to the different design of the regulation, which is tailored to encourage liquidity supply.

JEL Classification: G10, G20

Keywords: High Frequency Trading, Regulation, Order-to-Trade Ratio

1. Introduction

Stock exchanges are currently facing interesting times. In particular, the issue of designing the exchange's *trading rules* has become increasingly complex. When exchanges were mutual organizations, owned by its member stock brokers, they designed trading rules that would suit their members. Most of the world's exchanges have now demutualized into for-profit corporations. This has changed the exchanges main objective into maximizing profit for the exchange. To maximize revenues, exchanges now *compete* to attract order flow, where the design of trading rules has become a key tool.

This competition has also been enhanced by regulatory incentives such as Regulation National Market System (Reg NMS) in the United States, and MiFID in Europe, which explicitly enforce exchange competition. The new regulations have also introduced numerous competitors to the traditional exchanges through various forms of electronic OTC trading (dark pools) where the trading rules are much more opaque than the traditional exchanges that enforce pre-trade transparency.

The increased competition among exchanges has escalated the sophistication and diversity among the traders. The most important changes have come through the rise of electronic "algorithmic" traders (high-frequency traders–HFTs). The advent of HFTs has been met with scepticism by both regulators and the general public, in particular after the publication of *Flash Boys* by Michael Lewis (2014).

This has led to experimentation with the market design and trading rules, such as altering the calculation of trading

[☆]We would like to thank Thomas Borchgrevink at the Oslo Stock Exchange for providing us with information and data for the Oslo Stock Exchange. We also thank Christian Ringstad at the OSE for help with the data extraction. We are grateful for comments from a referee that helped us to significantly improve the paper. The paper has benefited from comments at conference presentations at the 2014 Central Bank Conference on Microstructure of Financial Markets in Italy, the 2015 FIBE conference, and the 2015 Arne Ryde Workshop at the University of Lund, the 2016 NYU Microstructure Meetings and the 2016 European Finance Association meetings in Oslo, as well as seminar presentations at the Universities of Göteborg and Stavanger. The views expressed are those of the authors and should not be interpreted as reflecting those of Norges Bank.

*Corresponding Author

¹BI Norwegian Business School, kjell.jorgensen@bi.no

²Norges Bank (Central Bank of Norway), jst@nbim.no

³UIS Business School, University of Stavanger(UiS), NO-4036 Stavanger, Norway, bernt.a.odegaard@uis.no

costs (make/take fees),⁴ changing tick sizes,⁵ payment for order flow,⁶ co-location,⁷ etc. In designing these trading rules, exchanges are balancing different profit components. As part of their efforts to attract order flow, their main source of revenue, exchanges want to make it easy for HFTs to access their limit order books. However, having HFTs actively placing orders in the order book might come with costs, both for the exchange, and for other traders accessing the same book. One is the IT costs, the need to deal with high-speed communications, and fast processing of the continuously updated limit order book. Another potential cost is reputational loss. Given the generally unfavorable view of HFTs, being too positive on HFTs may lead other traders to abandon an exchange. There is also the potential that regulators may introduce regulation not suited to an exchange's business model.

As a consequence, exchanges have tried to find ways to affect the incentives of HFTs that limit some of their message traffic, but that do not seriously inhibit their incentives to send order flow to the exchange. From the point of view of the exchange, and all other traders than the HFTs, the IT costs necessary to cater to the HFTs are economic externalities forced on them by the presence of HFTs. The economic problem for an exchange is similar to, for example, road congestion. According to economic theories of optimal taxation, the way to deal with such externalities is to design a tax that incentivizes behavior to avoid paying the tax. In the case of road congestion, the main goal of a tax is to encourage drivers to spread their driving away from rush hour, not primarily to generate revenues. The problem from an exchange's point of view is similar. The exchange's goal is to change the HFTs behavior by having them internalize the cost of excessive communications in their decision problems.

In this paper we look at a case where the Oslo Stock Exchange (OSE) in 2012 introduced a fee payable by traders with "excessive" order activity (order placements, order modifications, order withdrawals) relative to the number of trades in which the trader participates. The threshold that defined excessive order activity was set to an order-to-trade ratio (OTR) of 70:1 per month. That is, traders that post orders in excess of 70 per trade they participate in must pay a fee. However, the exchange realized that many of these orders are beneficial, as they provide liquidity for the market as whole, not merely to the fastest HFT.

To create incentives to leave an order in the book for some time, the exchange decided not to include orders staying in the book for more than one second when calculating the OTR. In addition, the OTR calculation excludes price-improving orders. The design of OSE's OTR fee thus has some of the same goals as a make-take fee structure; to encourage liquidity provision.

We investigate the consequences of the introduction of this fee on market quality at the OSE. We look at the impact on liquidity and trading volume on the exchange. We also compare the OSE to its closest competitors. The latter is relevant since although the fee was introduced at the OSE only, the OSE has a number of competing trading venues like NASDAQ-OMX (Stockholm), Chi-X, and BATS, where OSE-listed stocks are also traded. Traders may react in ways that the OSE desires, such as cutting down unnecessary message traffic while maintaining trading at the exchange. However, traders may also choose to move some or all of their trading activity away from OSE if they expect the OTR threshold to become binding, such that the cost of trading OSE-listed stocks is lower at other venues.

Our focus in this paper is changes in equity market structure and trader behavior, with a particular focus on the consequences of high-frequency trading on market quality. There is a large literature on this area, and we use the recent survey of Menkveld (2016) to summarize the main points.

The debate has centered around the contribution of HFTs to trading quality. Detractors argue that HFTs prey on all other traders in the market, and the playing field should be leveled.⁸ Others argue that much of what HFTs do is market making, the classical function of providing liquidity and being paid by earning the spread. This dichotomy led Menkveld and Zoican (2016) to classify HFTs into two types: "HFT bandits" and "HFT market makers." Much of the

⁴See U.S. Securities and Exchange Commission, Division of Trading and Markets (2015). For a summary of the arguments around make/take fees, see Angel, Harris, and Spatt (2011). For some empirical evidence on make/take fees, see Malinova and Park (2015) and Battalio, Corwin, and Jennings (2016).

⁵For an overview of the literature on tick sizes, see the survey by the U.S. Securities and Exchange Commission (2012). For an analysis of tick sizes on the OSE see Meling and Ødegaard (2017).

⁶See U.S. Securities and Exchange Commission, Division of Trading and Markets (2016).

⁷See Brogaard, Hagströmer, Nordén, and Riordan (2015) for an example study of co-location.

⁸A common suggestion is a change of trading protocols towards auctions (Budish, Cramton, and Shim, 2015). We also see market participants attempting to change protocols and trying to attract order flow. A well-known example is the IEX exchange, which has built in a delay in communications to and from the exchange in an attempt to deter certain kinds of HFT.

discussion around HFTs boils down to identifying the relative proportions of these two types.

Most of the theoretical literature on high-frequency trading is concerned with ways in which the HFTs can use their speed to get a competitive advantage,⁹ and is mainly suggesting that these activities have a negative impact on market quality. The findings in the empirical literature are more positive. First, there is a part of the literature that documents a decline in trading costs coinciding with the increased automation of trading.¹⁰ A second part of the literature documents that a substantial part of high-frequency trading activity is market making.¹¹

Exchanges are continuously adjusting their trading protocols. The open question is whether we have, in some sense, a socially optimal environment for trading financial assets. While any financial market place is unlikely to be a level playing field, just because of the fixed costs of accessing and analyzing it, there should be no obvious dead-weight losses to society,¹² and no sources of unfair advantages to the, for example, fastest traders.

The modern trading environment is a highly complex one. It has therefore become common to rely on studies of changes to market environments to seek out the more general relationships. Such empirical studies of the consequences of changes to exchange rules have the potential of putting a microscope on those aspects of the trading environment one is interested in. Our study is of this type. We look at the time period when the OSE introduced their OTR fee, and evaluate the effect on market quality. Our results suggest that there were no *negative* effects on liquidity and other measures of trading quality at the OSE.

These results are interesting for several audiences. They are directly relevant for the exchanges themselves, in addition to market regulators, and points to an OTR as a viable way of regulating order activity. The results also have more general implications for understanding HFTs, as we find some evidence that they changed their behavior around the introduction of the fee. This may also be what one would expect, since we know that HFTs are very cost sensitive, and successful HFT algorithms would deal with a potential fee by building the rules for the fee into their algorithm.

Our results are also interesting because they give a different message than two other studies of similar introductions of OTR fees in other markets. For example, Friederich and Payne (2015) study the introduction of an OTR fee in the Italian market, where a fee is paid on OTRs in excess of 100:1. The Italian fee on the exchange was politically motivated, which may explain why the Italian fee is different from the one at the OSE. It is calculated on a daily basis, and there are no exceptions for long-lived orders or price-improving orders. Therefore, the fee has no associated incentives for liquidity provision as opposed to the fee introduced at the OSE. This may explain the negative effects on liquidity they document. They find a 15% increase in spreads in the market at the time of the introduction of the fee.

Another study is Malinova, Park, and Riordan (2016), who look at evidence from the Canadian market. Canada introduced a fee with similar implications as an OTR. Canadian traders are paying the Canadian regulator IROC a fee for surveillance services. This fee was initially calculated on the basis of the number of trades. The fee calculation was changed in 2012 to be based on both trades and orders (messages) into the limit order book. At the time of the introduction of the measure, it was unclear how the calculation actually was to be done, with corresponding uncertainty from market participants about the fee they would end up paying at the end of the month. Malinova et al. show that the measure led to an immediate 9% increase in quoted bid-ask spreads, with corresponding cost increases for retail traders. They show that this spread increase is caused by HFT market makers pulling back. The difference between the OSE and these other markets shows that regulatory design matters.

Our study also relates to other studies that examine the effect of direct market interventions from regulators on the behavior of HFTs. An example is the event where the U.S. Securities and Exchange Commission (SEC) enforced the rule that all trades have to go through exchange members, without possibilities for traders to directly enter orders into the exchange's limit order books (the Naked Access Ban). This is a direct intervention aimed at slowing down HFTs.¹³ Another example is introductions of co-location, which are occasions where HFTs are enabled to speed up. So, when looking at HFTs in general, our results are related to these other studies of changes in HFT functionality.

To summarize our main findings, we find no negative effects on market quality at the OSE in response to the introduction of the OTR fee. We support this finding with a number of complementary analyses. We first examine

⁹See e.g. Biais, Foucault, and Moinas (2015) and Foucault, Hombert, and Roşu (2016).

¹⁰Angel et al. (2011) shows the trends, Hendershott, Jones, and Menkveld (2011) links it more specifically to automated trading.

¹¹See for example Hagströmer and Nordén (2013) and Menkveld (2013).

¹²An example often pointed to is the huge costs of shaving a microsecond from the communications lag between New York City and Chicago (Laughlin, Aguirre, and Grundfest, 2014; Lewis, 2014).

¹³See e.g. Chakrabarty, Jain, Shkilko, and Sokolov (2014).

whether the OTR limit of 70:1 was likely to be one that most HFTs would view as binding. We show that, at the aggregate level, there were days with individual stocks with a market-wide OTR higher than 70:1. We investigate the individual exchange members and calculate monthly OTRs for each member and each stock. Our results show that in May 2012, the month the OSE OTR fee scheme was announced, there were some extreme cases of members with OTRs above 1000:1 for a stock. Hence, the limit of 70:1 was clearly binding for some participants.

When we look at the trading decisions of individual members, we see that they do react to the regulation. Members with high OTRs in a given stock in May 2012 have much lower OTRs in the same stock in September 2012, which is the month the scheme went into effect. We also show that members with high OTRs in the first half of September 2012 reduce their OTRs for the second half of that month, again on a stock by stock basis.

Our main findings are based on the evolution of various measures of liquidity. We complement the liquidity measures by a comparison of *price discovery*, by estimating the information share of the OSE and its main competitor. We investigate whether there are changes to the relative importance of OSE in price discovery linked to the introduction of the fee. We find no such effects, the OSE's share of price discovery seems largely unchanged around the event.

The remainder of the paper is structured as follows. In Section 2, we give some background on the market place and data sources. In Section 3, we present the details on the fee introduction. Our main results are given in Section 4, where we show what happens to market quality at the OSE when the fee was introduced. We split this investigation into several parts, first showing some longer term trends around the introduction of the fee, before looking more directly at changes just in the months of the introduction through a series of diff-in-diff analyses. The main analysis is then complemented in Section 5, before we conclude in Section 6.

2. Market place and data

In this section we present the Oslo Stock Exchange and other venues for trading in Norwegian equities, give our data sources, and some descriptive statistics.

2.1. The Oslo Stock Exchange

Norway is a member of the European Economic Area (EEA) and its equity market is among the 30 largest world equity markets by market capitalization. The OSE is the only regulated marketplace for securities trading in Norway. Unlike the other Scandinavian exchanges, the OSE has remained relatively independent, but has been in strategic partnership with the London Stock Exchange (LSE) since March 2009.

Since January 1999 the OSE has operated as a fully computerized limit order book.¹⁴ As is normal in most electronic order-driven markets, the order handling rule follows a strict price-time priority. All orders are submitted at prices constrained by the minimum tick size.¹⁵ The trading day at the OSE comprises three sessions: an opening call period, a continuous trading period, and a closing call period. There may also be call auctions in the continuous trading period for any security if triggered by price monitoring, or to restart trading after a trading halt. The orders are matched in accordance with their priority, which is price-visibility-time for round-lot orders.¹⁶ In September 2012, the continuous trading session was changed from 09:00 to 17:20 to 09:00 to 16:20.

The distribution of firm size and trading volume at the OSE is heavily skewed. The OSE is dominated by a few very large companies, of which the largest, Statoil, an oil company, at the beginning of 2012 accounted for about 25% of OSE market capitalization. Two other companies, Telenor (telecommunications) and Den Norske Bank (integrated financial) each accounted for about 10% of OSE market capitalization. The large firms at the OSE dominate the trading volume at the exchange. Trading interest is concentrated in the constituents of the OBX index, which contains the 25 most liquid stocks at the OSE.¹⁷

¹⁴For further background on the trading at the OSE and the companies on the exchange, see Bøhren and Ødegaard (2001), Næs and Skjeltorp (2006), and Næs, Skjeltorp, and Ødegaard (2011).

¹⁵For details about tick sizes, see Meling and Ødegaard (2017).

¹⁶With the OSE's migration to TradElect in partnership with the London Stock Exchange Group in April 2010 the OSE offers its members the opportunity to preferentially trade with themselves before trading with other participants when there is more than one order at a given price level. This means that orders submitted for a trader configured to use Own Order Preferecing will execute in the following order: Price-Counterparty-Visibility-Time.

¹⁷See Meling (2016) for more details on the OBX index.

2.2. Migration of trading to alternative market places

Post MiFID, the trading of stocks with a main listing at the OSE has become increasingly fragmented across various alternative market places. In our sample period the largest European competitors with pre-trade transparency (i.e. limit order books), are the Stockholm Stock Exchange (Nasdaq OMX Nordic), Chi-X, BATS, and Turquoise. Due to the strategic partnership with LSE, there is little direct competition between LSE and OSE in stocks with a main listing at the OSE. Some of the largest stocks at the OSE are also traded overseas, such as at the NYSE and NASDAQ. In this paper, we do not consider trading outside of Europe, as this tends to be in other time-zones, with little overlap in opening hours.¹⁸

In addition to these market places with pre-trade transparency, there are also numerous alternative market places facilitating OTC or dark pool trading, where transparency is only ex post. These market places are required to report their trades to a MiFID-compliant reporting facility. We use all trades reported through a major reporting facility, Markit BOAT, to proxy for OTC trading in our sample of OSE listed stocks.

Not all stocks listed at the OSE are traded elsewhere. Only the larger companies on the exchange are interesting for the competing market places. The OSE lists between 200 and 300 stocks, of which only about 50 have a significant amount of trading outside the OSE.

2.3. Data sources

We rely on a number of datasets to analyze the trading in stocks with a main listing at the OSE. First, we use a dataset from the order book at the OSE provided by the market surveillance department at the exchange. This dataset provides information about all trades and orders at the exchange. The dataset also includes various additional information about each order, such as order cancellations, order modifications (volume and/or price updates), hidden orders, etc., which makes it possible to construct OTRs. The data also includes an anonymized identifier linking stock exchange members to each order and trade.

We also use the Thomson Reuters Tick History Database, which contains information for all European market places where stocks with a main listing at the OSE are traded. While this dataset also contains orders, trades, and the state of the order book, there is less additional information compared to the OSE data. For lit market places, markets with pre-trade transparency, the dataset includes the ten best levels of the bid and ask side of the limit order book. However, the data do not allow us to construct approximations of OTRs, as there is not a complete record of order messages to the different exchanges. The Thomson Reuters data also include some information about OTC trading of OSE stocks, through the inclusion of trades reported through Markit BOAT.

Finally, we have data from the OSE Information Service (OBI), which provides daily price observations, together with information about corporate events, corporate announcements, and accounts.

In the main analysis, we use data for equities with a main listing at the OSE. We only use common equity and exclude ETFs and other equity-like instruments. In 2012, there was a total of 243 equities listed. We remove the least liquid stocks by only including stocks with a minimum of 100 trading days in a year, which reduces the sample to 119 stocks.

2.4. Market quality measures

We estimate a number of standard empirical measures of market quality. Market *Depth* is calculated as the sum of trading interest at the best bid and ask, in Norwegian kroner (NOK). In the analysis, we use the daily average of the depth each time there is an update of the order book.

A number of spread measures are calculated using the full trading record. The *relative spread* is the difference between the current best ask and best bid, divided by the average of these. We use all events with an update of the state of the limit order book. For each update, we calculate the relative spread using the currently best bid and best offer. Our estimate of that day's relative spread is the sample average over the day.

The *Effective Spread* relates transaction prices to the spread when the order is submitted. We calculate the effective proportional spread as $q_{jt}(p_{jt} - m_{jt})/m_{jt}$, where q_{jt} is an indicator variable that equals +1 for buyer-initiated trades and

¹⁸Between 2009 and September of 2012 there was an overlap of one hour between trading in Oslo and New York City, which disappeared when Oslo moved their closing time from 17:20 to 16:20 local time. In our empirical work, we only look at trading when the OSE is open.

-1 for seller-initiated trades, p_{jt} is the trade price, and m_{jt} is the quote midpoint prevailing at the time of the trade. To determine whether an order is buyer or seller initiated, we compare the price to the midpoint. If the price is above the midpoint, we classify it as buyer initiated. Otherwise, we classify it as seller initiated. In the analysis, we use the daily average of effective spreads for all trades during the day.

The *Realized Spread* is calculated as $q_{jt}(p_{jt} - m_{j,t+5\text{min}})/m_{jt}$, where p_{jt} is the trade price, q_{jt} is the same buy/sell indicator as that used for the effective spread, m_{jt} is the prevailing midpoint, and $m_{j,t+5\text{min}}$ is the quote midpoint 5 minutes after the t 'th trade. Similarly to the effective spread, we calculate the daily average of realized spreads for all trades during the day.

To measure the variability of prices, we use the *Realized Volatility*, estimated as the second (uncentered) sample moment of the return process over a fixed interval of 10 minutes, scaled by the number of observations n . We calculate the realized volatility on a daily basis.

Finally, we calculate a rougher measure of trading costs, the *Roll* measure. This is an estimate of trading cost that uses the autocovariance induced by bid/ask bounce to estimate the size of the implicit spread between bid and ask prices. Our motivation for the inclusion of the Roll measure is that it can be calculated in situations where we do not have an order book, just prices. This is the case for the OTC data, which are only reported post-trade. We calculate the Roll measure based on the returns r_t calculated from transaction-to-transaction prices during a day. The Roll spread estimator is $\hat{s} = 2\sqrt{-\text{cov}(r_t, r_{t+1})}$. We only use observations where the autocovariance is negative.

In Table 1, we describe these measures using data for the period before the introduction of the OTR fee, 2010–2011. We report averages of daily estimates, both for the whole market, and for size-sorted portfolios.

Table 1 Descriptive Statistics for the Market Quality Measures: 2010-2011

	Whole	Size Portfolio			
	sample	1 (small)	2	3	4
Depth (thousands NOK)	279	104	176	326	547
Relative Spread (%)	2.95	4.94	4.39	1.99	0.81
Effective Spread (%)	0.87	1.86	1.02	0.51	0.25
Realized Spread (%)	0.32	0.73	0.31	0.20	0.10
Realized volatility (%)	0.64	1.10	0.68	0.47	0.39
Roll (%)	0.47	1.05	0.45	0.30	0.12

We describe the measures of market quality used in the analysis. The numbers in the tables are averages of daily estimates. Depth is the sum of trading interest (in thousands NOK) at the best bid and best ask. The relative spread is the difference between best bid and best ask scaled by the prevailing midpoint. The effective spread is transaction prices minus the prevailing midpoint just before the transaction. The realized spread is the transaction price minus the prevailing midpoint five minutes after the transaction. Both realized and effective spreads are multiplied with a trade direction indicator, and scaled by the prevailing midpoint. Daily depth and spread measures are calculated as averages across intraday observations. The realized volatility is the (uncentered) second moment of 10-minute returns. The Roll measure is calculated from the autocovariance of trade to trade returns. The sample excludes illiquid stocks, which are stocks with less than 100 trading days in a year.

3. Introduction of the order-to-trade ratio at the OSE

In terms of terminology, the OSE uses the term “Order to Executed Order Ratio” for their specific definition of an OTR, which accounts for the duration of an order and whether it is price improving.

3.1. Specifics of the OSE order to executed order ratios (OEOR) fee

The introduction of the fee on excessive Order to Executed Order Ratios (OEOR) was announced by the OSE on May 25, 2012. The announcement justified the introduction on efficiency grounds, arguing that excessive order activity was imposing negative externalities on all market participants. The full text of the press release is given in Figure 1.

In addition to the press release, the OSE also gave more details about the actual fee structure and the calculation of the OTR.¹⁹ The calculation is done on a monthly basis. The actual fee is NOK 0.05 per message that exceeds a ratio of

¹⁹“Oslo Børs to Implement Order to Executed Ratio,” downloadable from the OSE website (oslobors.no).

Figure 1 Press Release, May 25, 2012, from the Oslo Stock Exchange

With effect from 1 September, Oslo Børs will introduce a fee that will affect unnecessarily high order activity in the stock market. The purpose of the fee is to discourage orders that do not contribute to the effective and sound conduct of stock market trading. Order activity at unnecessarily high levels has the effect of reducing the transparency of the order picture and so reducing confidence in the market.

Competition and technological development have played a role in radical changes in trading behaviour in the stock market over recent years. Increased use of algorithms as a tool for carrying out various kinds of trading strategy has resulted over time in a steady reduction in the average order size, combined with an increase in the number of order events relative to the number of trades actually carried out. This creates both direct and indirect costs for all market participants, due in part to greater volumes of data and the requirements this creates in terms of investment in infrastructure and greater bandwidth.

“Oslo Børs takes the view that high order activity is not in itself necessarily negative for the market, but we are keen to encourage a situation in which all types of trading contribute to maintaining confidence in the marketplace,” comments Bente A. Landsnes, President and CEO of Oslo Børs.

“It is in general the case that a market participant does not incur any costs by inputting a disproportionately high number of orders to the order book, but this type of activity does cause indirect costs that the whole market has to bear. The measure we are announcing will help to reduce unnecessary order activity that does not contribute to improving market quality. This will make the market more efficient, to the benefit of all its participants,” explains Bente A. Landsnes.

The fee will be linked to an “Order to Executed Order Ratio (OEOR)” of 1:70. This means that the fee will be charged where the number of orders input relative to each order carried out exceeds 70. The order activity that will be included in the calculation of this ratio will principally relate to orders that are cancelled or amended within one second, and where the change does not contribute to improved pricing or volume.

Accordingly, orders that remain open in the order book for some time, or which are updated in a manner that makes a positive contribution to market quality by reducing the spread between best bid and best offer or by increasing order book depth will not be included in the calculation of the type of activity that Oslo Børs wishes to make the subject of the additional fee.

70:1. In the calculation, the OSE does not count every message. Specifically, orders with the following characteristics are excluded from the calculation:

- Orders that rest unchanged for more than one second from entry.
- Order amendments that improve price, volume, or both.
- Execute and Eliminate (ENE) and Fill or Kill (FOK) orders.

Orders that have the following features *are* counted:

- Orders residing less than one second, from order insert or the last amendment, before cancellation.
- Order amendments that degrade price, volume, or both, of an order that has resided for less than one second in the trading system.

The way executed orders are counted is also specified as:

- Orders that result in one or many transactions are counted as one executed order.
- Executed orders, orders that have been involved in one or more trades, but with total executed value of less than NOK 500, will not be counted as an executed order.

3.2. Possible Trader Reactions

There are many possible ways traders can react to the OTR fee. One is to ignore the new potential fee. Another option is to switch trading from the OSE to one of the other exchanges trading in OSE listed stocks. For HFTs who wish to stay at the OSE, the more natural reaction is to reprogram their algorithms, building in the rules of the OTR calculation, and factoring in the potential cost in the algorithm's actions. There are a set of potential changes that traders could implement. One possibility is to change the size of the orders. If more trades are needed, one could lower order sizes. But there is a limit to how small the orders can be given that the fee has a built in lower limit of NOK 500 before a trade counts. So, for example, for a stock priced at 50, the trade needs to be higher than 10 shares. Traders that split large orders into smaller pieces could consider increasing the size of the individual pieces. Malinova et al. (2016) show that in a similar setting, market makers react by widening the spread, presumably to reduce the option value of longer-lived limit orders.

The choice of reaction will depend on the trading strategy and type of trader. Many traders, in particularly buy-side traders, which are naturally long in equities, are not likely to run into the OTR threshold of 70:1 when buying and selling equities for long-term portfolio purposes, even if they are using an order-splitting algorithm.

The traders who are more likely to have the fee "bite" are HFTs of some kind. Their reactions depend on the type of strategy a given trader is following. As already mentioned, the OSE's regulatory change is designed to be less onerous for market making strategies. To see why, recall that market making involves placing orders to buy and sell in the limit order book, hoping to earn the spread. When prices change, these orders are updated. If the market maker maintains the spread and updates her limit orders to the new price level, either the bid and ask will be price improving and hence not counted in the OTR calculation. In other words, for market makers maintaining the same spread, only *half* of the new orders will count in the calculation of the OTR. Similarly, when there is little market activity in the market, market makers' quotes are likely to stay in the order book for longer than one second and will therefore also not count in the calculation of the OTR. It is apparent that the design of the fee calculation rewards liquidity provision.

An high-frequency trading strategy that is more likely to get a high OTR is a "relative value strategy." Here, a trader reacts to price discrepancies between two or more market places. The strategy involves sending orders to both exchanges at current prices, orders that needs to be filled immediately. Such orders are neither price improving nor long lived and will all count in the calculation of the OTR.

HFTs that run on public news to pick off the quotes of others, which is one rumored strategy of HFT "bandits," would also have all their orders count. Other hypothesized strategies of the HFT bandits, such as "spoofing" (posting large orders outside the spread), "smoking" (posting fleeting orders inside the current spread), and "stuffing" (posting

many orders slowing down communications to the exchange for other traders) would also risk hitting the OTR threshold and be charged the fee, as these strategies involve orders that are not meant to be executed. This kind of behavior is something that exchanges want to actively discourage.

We also investigate some aspects of the reaction functions of individual traders. We examine whether individual traders that are likely to be affected by the fee react in order to lower their OTR. For example, a trader with a high OTR in the beginning of the month may consider “pushing back” trading in that stock towards the end of the month, to avoid paying the fee.

In their definition of the specific OTR calculated by the OSE, the OSE uses the term “OEOR” for the calculation done for each trader, accounting for whether orders are long lived or price improving (and hence not counted), order size, and various other factors, such as a minimum trade value. We do not attempt to approximate such a calculation. Instead, we calculate OTRs using all orders and all trades, either for all traders in a given stock (market aggregates), or for a single trader in a given stock. The OTRs we estimate may therefore be slightly higher than the actual OEOR in the exchange’s system, but using a generic OTR facilitates comparisons with other exchanges, and also with the period before the OEOR was introduced.

4. Market Quality Changes

We first investigate how liquidity was affected in a broad sense around the introduction of the OTR fee, before looking in more detail at the month when the scheme was announced (May 2012) and the first month in which the fee was payable (September 2012). We perform basic pre and post comparisons to investigate how broad measures of liquidity changed over the event. As with any such comparison across time, there may be other confounding factors. We therefore also look at difference in difference specifications with better econometric properties. We first look at the picture on the OSE itself, where we compare the high OTR stocks, those that were likely to be most affected by the OTR introduction, with stocks unlikely to be affected by the constraint of 70:1 (i.e. firms with low OTRs). We then look at the possibilities of liquidity moving away from the OSE by comparing liquidity at the OSE with the liquidity at the largest non-OSE market for a particular stock.

4.1. Changes in liquidity

The most important question is whether market quality is affected by the introduction of the OTR fee. We first calculate market quality measures pre and post the introduction of the OTR fee. The fee payable for traders with an OTR above 70:1 was introduced on September 1, 2012. We use the period September-November 2012 (“Fall 2012”) to measure behavior *post* introduction. This is compared to two alternative “pre” periods: The corresponding period (September-November) the year before (“Fall 2011”) as well as the period from the beginning of 2012 until the announcement of the OTR fee in May of 2012 (“Spring 2012”). If there are seasonalities in trading, the cleanest comparison will be the previous fall. The spring is however closer in time. We therefore examine both.

We first look at the OSE in isolation. Table 2 shows averages of market quality measures for the whole market. With the exception of the relative spread, all of the quality measures are significantly improved relative to both the previous spring and fall. Table 2 also reports averages of the OTR for the same periods. The average OTR is reduced, especially compared to the previous fall.

We also examine these statistics across different market capitalization (size) groups. Table 3 shows similar numbers to those in Table 2 for the four size-sorted portfolios. Here, we see much of the same. Most of the size-based portfolios also show significant improvements in liquidity, especially when we compare Fall 2012 to Fall 2011. There are a few cases, however, where the liquidity deteriorates. This happens to the smallest companies for the depth and relative spread measures when we compare Spring 2012 to Fall 2012.

If market quality deteriorates for stocks where the OTR is binding, we can attempt to identify this by looking separately at the stocks for which it is expected to be binding. We therefore take a close look at stocks where the OTR exceeded the limit of 70:1 at least once during 2011. Table 4 shows that from Fall 2011 to Fall 2012, all market quality measures show significant improvements for the OTR > 70 group. For the comparison between Spring 2012 and Fall 2012, the picture is less unanimous, but all significant estimates show a quality improvement for the groups with max OTR > 70, those most likely to be constrained by the limit.

Table 2 Trade quality measures before and after the OTR fee introduction

	Averages			Test for equality(<i>p</i> -value)	
	Fall 2011	Spring 2012	Fall 2012	Fall 2011 vs Fall 2012	Spring 2012 vs Fall 2012
Order-to-Trade Ratio	24.4	17.3	16.7	-11.5 (0.00)	-1.4 (0.17)
Depth (thousands NOK)	200	267	291	19.7 (0.00)	4.8 (0.00)
Relative Spread (%)	3.43	2.84	2.82	-24.1 (0.00)	-0.8 (0.42)
Effective Spread (%)	0.88	0.66	0.59	-25.2 (0.00)	-6.9 (0.00)
Realized Spread (%)	0.30	0.23	0.19	-11.7 (0.00)	-5.5 (0.00)
Realized volatility (%)	0.81	0.60	0.51	-21.3 (0.00)	-8.3 (0.00)
Roll (%)	0.54	0.42	0.36	-10.2 (0.00)	-4.4 (0.00)

The table provides results for measures of liquidity for three subperiods: Fall 2011: September–November 2011; Spring 2012: January–May 2012; Fall 2012: September–November 2012. For each measure, we calculate it on a daily basis for all stocks in the sample. The reported numbers are averages of these daily estimates. The OTR is the number of orders (messages to the limit order book) divided by the number of trades. Depth is the sum of trading interest (in thousands NOK) at the best bid and best ask. The relative spread is the difference between best bid and best ask scaled by the prevailing midpoint. The effective spread is transaction prices minus the prevailing midpoint just before the transaction. The realized spread is the transaction price minus the prevailing midpoint five minutes after the transaction. Both realized and effective spreads are multiplied with a trade direction indicator, and scaled by the prevailing midpoint. Daily depth and spread measures are calculated as averages across intraday observations. The realized volatility is the (uncentered) second moment of 10-minute returns. The Roll measure is calculated from the autocovariance of trade to trade returns.

Tables 3 and 4 also show the development of the OTR for these different groups. Note that for large companies, and stocks with high OTR, the OTR has fallen. Thus, by looking at the OSE in isolation, there seems to be little evidence of negative effects on market quality coinciding with the introduction of the OTR fee.

We also look at potential effects in other market places where OSE-listed companies are traded by attempting to measure *aggregate* market quality for these stocks across all lit markets. To do this we use Reuter’s aggregate summary of trading across Europe captured by their XBO feed. This is a record of time-stamped trades at all European market places. However, it is only constructed for stocks with significant cross-exchange trading. Hence, it is only a subset of all stocks (the largest) used in the previous estimates from the OSE. Furthermore, the feed only includes prices and volumes, which means that we are limited to measures that can be estimated from trade sequences. In Table 5, we show the averages of two such market quality indicators: realized volatility and the Roll implicit spread estimator. The realized volatility fell significantly after the introduction of the fee while the Roll measure seems relatively unchanged, as the estimate for the Fall 2012 is between the previous fall and spring. Again, we see no evidence of reduced market quality following the introduction of the fee.

4.2. Did market quality change at the OSE?

A shortcoming of the difference analysis above is that it is not obvious that the changes we see are truly related to the OTR fee introduction, as they can be driven by longer term trends. We therefore perform several analyses to more definitively identify any effects related to the fee introduction.

We first use the fact that the fee introduction does not affect all the stocks on the OSE equally. Many of the smaller stocks on the OSE do not have such high OTRs that the introduction of the fee is likely to affect them. We can therefore use these “low OTR” stocks as an untreated sample in a difference in differences analysis. Figure 2 illustrates the analysis. The fee was announced at the end of May 2012, and implemented starting September 2012. We use the first part of the year (January–April) to choose a set of stocks not likely to be affected by the new regulation. To proxy for that, we first measure the OTR for each stock for each day over the January–April period. We presume that stocks with a maximal OTR of less than 50 over this period are not likely to be affected by the introduction of the fee. This group of stocks constitute our control sample. We term this group the “Low OTR” stocks. This is then compared to a treatment sample; stocks with an observed OTR higher than 50 in the same time period. We term these the “High OTR” stocks.

Estimation of the difference in differences is based on regressions of the following type:

$$y = \beta_0 + \beta_1 d_{treated} + \beta_2 d_{time} + \delta d_{treated} \times d_{time} + \alpha \mathbf{X} + \varepsilon, \quad (1)$$

where y is the variable of interest (i.e., a liquidity measure), $d_{treated}$ is a dummy variable for whether an element belongs to the treatment or the control group (high vs. low OTR), and d_{time} a time dummy for the second period. The coefficient

Table 3 Trade quality measures before and after the OTR fee introduction: size-sorted portfolios

	Size Quartile	Averages			Test for equality(p -value)	
		Fall 2011	Spring 2012	Fall 2012	Fall 2011 vs. Fall 2012	Spring 2012 vs. Fall 2012
Order-to-Trade Ratio	1 (small)	10.7	10.8	11.1	1.9 (0.06)	1.4 (0.15)
	2	13.8	11.8	15.8	3.3 (0.00)	7.3 (0.00)
	3	55.6	32.8	25.8	-10.6 (0.00)	-4.7 (0.00)
	4	48.6	27.8	20.3	-9.7 (0.00)	-9.1 (0.00)
Depth (thousands NOK)	1 (small)	80	124	115	11.7 (0.00)	-2.2 (0.03)
	2	175	197	226	5.8 (0.00)	3.5 (0.00)
	3	220	322	357	12.6 (0.00)	3.0 (0.00)
	4	354	484	501	13.4 (0.00)	1.4 (0.15)
Relative Spread (%)	1 (small)	5.82	5.40	5.64	-3.0 (0.00)	3.5 (0.00)
	2	5.13	3.82	3.72	-22.6 (0.00)	-1.5 (0.13)
	3	2.06	1.77	1.71	-12.0 (0.00)	-1.8 (0.07)
	4	0.91	0.69	0.65	-17.5 (0.00)	-3.4 (0.00)
Effective Spread (%)	1 (small)	1.83	1.46	1.32	-20.0 (0.00)	-5.5 (0.00)
	2	1.03	0.69	0.55	-23.6 (0.00)	-7.4 (0.00)
	3	0.52	0.43	0.43	-8.1 (0.00)	-0.2 (0.82)
	4	0.20	0.16	0.15	-8.0 (0.00)	-1.5 (0.13)
Realized Spread (%)	1 (small)	0.82	0.61	0.46	-11.7 (0.00)	-6.1 (0.00)
	2	0.27	0.20	0.17	-8.0 (0.00)	-3.6 (0.00)
	3	0.14	0.14	0.12	-2.4 (0.02)	-2.2 (0.02)
	4	0.04	0.04	0.03	-1.6 (0.10)	-2.3 (0.02)
Realized volatility (%)	1 (small)	1.39	1.16	0.99	-12.7 (0.00)	-6.0 (0.00)
	2	0.87	0.59	0.47	-23.2 (0.00)	-8.3 (0.00)
	3	0.54	0.44	0.39	-13.0 (0.00)	-5.1 (0.00)
	4	0.50	0.33	0.28	-6.4 (0.00)	-1.7 (0.09)
Roll (%)	1 (small)	1.36	1.01	0.77	-11.3 (0.00)	-5.8 (0.00)
	2	0.51	0.38	0.35	-10.2 (0.00)	-2.5 (0.01)
	3	0.29	0.27	0.25	-3.0 (0.00)	-1.2 (0.25)
	4	0.12	0.09	0.10	-0.7 (0.49)	0.3 (0.79)

The table provides results for properties of liquidity for three subperiods: Fall 2011: September-November 2011; Spring 2012: January-May 2012; Fall 2012: September-November 2012. For each measure, we calculate it on a daily basis for all stocks in the sample. The reported numbers are averages of these daily estimates for four size sorted portfolios. The OTR is the number of orders (messages to the limit order book) divided by the number of trades. Depth is the sum of trading interest (in thousands NOK) at the best bid and best ask. The relative spread is the difference between best bid and best ask scaled by the prevailing midpoint. The effective spread is transaction prices minus the prevailing midpoint just before the transaction. The realized spread is the transaction price minus the prevailing midpoint five minutes after the transaction. Both realized and effective spreads are multiplied with a trade direction indicator, and scaled by the prevailing midpoint. Daily depth and spread measures are calculated as averages across intraday observations. The realized volatility is the (uncentered) second moment of 10-minute returns. The Roll measure is calculated from the autocovariance of trade to trade returns. The sample excludes illiquid stocks, which are stocks with less than 100 trading days in a year.

Table 4 Trade quality measures before and after regulation: OTR sorted portfolios

	Max OTR 2011	Averages			Test for equality(p -value)	
		Fall 2011	Spring 2012	Fall 2012	Fall 2011 vs. Fall 2012	Spring 2012 vs. Fall 2012
Order-to-Trade Ratio	OTR<50	12.0	11.2	11.8	-0.9 (0.36)	3.2 (0.00)
	OTR \in [50,70]	12.0	12.0	12.0	-0.3 (0.79)	-0.0 (0.99)
	OTR>70	47.4	28.3	23.6	-12.9 (0.00)	-5.2 (0.00)
Depth (thousands NOK)	OTR<50	202	272	321	15.5 (0.00)	6.2 (0.00)
	OTR \in [50,70]	133	206	206	11.5 (0.00)	-0.0 (0.98)
	OTR>70	218	288	279	10.2 (0.00)	-1.5 (0.13)
Relative Spread (%)	OTR<50	4.14	3.50	3.54	-14.0 (0.00)	0.9 (0.39)
	OTR \in [50,70]	5.01	3.90	4.16	-10.1 (0.00)	3.0 (0.00)
	OTR>70	2.36	1.87	1.74	-21.3 (0.00)	-4.3 (0.00)
Effective Spread (%)	OTR<50	0.97	0.80	0.63	-17.7 (0.00)	-10.3 (0.00)
	OTR \in [50,70]	1.38	0.94	0.88	-16.5 (0.00)	-2.4 (0.02)
	OTR>70	0.70	0.51	0.49	-21.2 (0.00)	-1.7 (0.09)
Realized Spread (%)	OTR<50	0.28	0.27	0.17	-6.2 (0.00)	-7.4 (0.00)
	OTR \in [50,70]	0.54	0.33	0.27	-8.6 (0.00)	-2.8 (0.01)
	OTR>70	0.24	0.18	0.17	-8.1 (0.00)	-0.9 (0.36)
Realized volatility (%)	OTR<50	0.87	0.68	0.51	-16.0 (0.00)	-11.7 (0.00)
	OTR \in [50,70]	1.21	0.81	0.75	-12.2 (0.00)	-2.1 (0.04)
	OTR>70	0.67	0.50	0.45	-17.6 (0.00)	-4.8 (0.00)
Roll (%)	OTR<50	0.42	0.43	0.28	-5.9 (0.00)	-7.2 (0.00)
	OTR \in [50,70]	0.92	0.66	0.48	-7.2 (0.00)	-4.2 (0.00)
	OTR>70	0.51	0.36	0.34	-9.4 (0.00)	-0.8 (0.45)

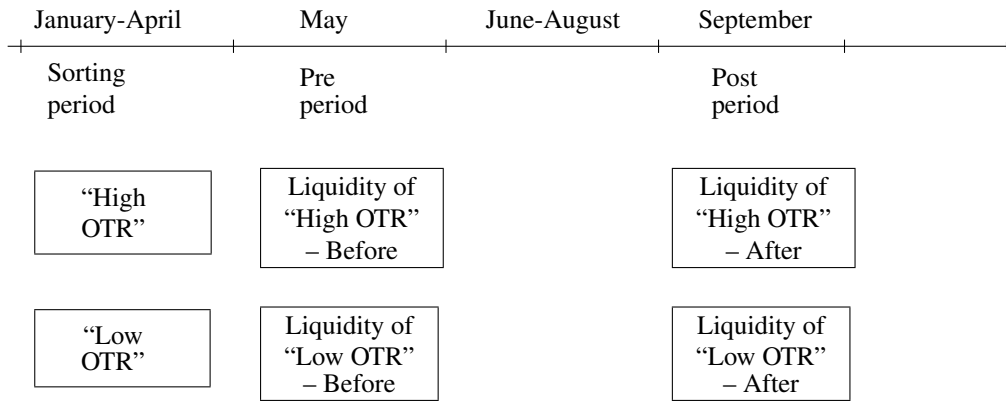
The table provides results for properties of liquidity for three subperiods: Fall 2011: September-November 2011; Spring 2012: January-May 2012; Fall 2012: September-November 2012. For each measure, we calculate it on a daily basis for all stocks in the sample. The reported numbers are averages of these daily estimates for three different groups: Stocks with max OTR<50, max OTR \in [50, 70], and max OTR>70. The OTR is the number of orders (messages to the limit order book) divided by the number of trades. Depth is the sum of trading interest (in thousands NOK) at the best bid and best ask. The relative spread is the difference between best bid and best ask scaled by the prevailing midpoint. The effective spread is transaction prices minus the prevailing midpoint just before the transaction. The realized spread is the transaction price minus the prevailing midpoint five minutes after the transaction. Both realized and effective spreads are multiplied with a trade direction indicator, and scaled by the prevailing midpoint. Daily depth and spread measures are calculated as averages across intraday observations. The realized volatility is the (uncentered) second moment of 10-minute returns. The Roll measure is calculated from the autocovariance of trade to trade returns. The sample excludes illiquid stocks, which are stocks with less than 100 trading days in a year.

Table 5 Aggregate European trading and liquidity before and after OTR fee introduction

	Averages			Test for equality(p -value)	
	Fall 2011	Spring 2012	Fall 2012	Fall 2011 vs. Fall 2012	Spring 2012 vs. Fall 2012
Realized volatility (%)	0.58	0.40	0.35	-20.7 (0.00)	-6.1 (0.00)
Roll (%)	0.12	0.10	0.11	-3.5 (0.00)	3.8 (0.00)

The table shows averages of market quality measures across European exchanges. For each stock where Reuters provide an aggregate (XBO), we calculate the market quality measures on a daily basis. The table reports averages across three subperiods: Fall 2011: September-November 2011; Spring 2012: January-May 2012; Fall 2012: September-November 2012. The Roll measure is calculated from the autocovariance of trade to trade returns. The realized volatility is the (uncentered) second moment of 10-minute returns.

Figure 2 Illustration of the difference in differences analysis



of interest, δ , multiplies the interaction term ($d_{treated} \times d_{time}$). This term also has the interpretation as a dummy variable equal to one for observations in the treatment group in the second period. The coefficient δ measures the direct effect of the intervention. We adjust for the panel data nature of the data by including fixed date and stock effects, and adjusting the standard errors in the panel for clustering.²⁰

In the regression in equation (1), we allow for additional covariates \mathbf{X} . We estimate two versions of the model with different additional covariates. In panel A of Table 6 we only control for size differences between the high and low OTR groups by including log firm size as an additional explanatory variable. In Panel B of Table 6, in addition to size, we also include the *Realized Volatility (RV)* and the inverse of the stock price ($1/Price$) as covariates. We also examine a regression with *RV* as dependent variable and report the results in Panel A. This specification is left out in Panel B, since we there use *RV* as an explanatory variable.

In addition to the liquidity variables, the first column in both panels of Table 6 includes the results of an estimation with the OTR as a dependent variable. This information can be used to investigate whether the reactions of traders on the OSE actually leads to relatively lower OTRs for the stocks with high OTRs relative to low OTRs. The coefficient on the interaction term is -11.3 (-10.9 in the specification with additional covariates), which suggests that the OTR of the stocks that are likely to hit the OTR limit reduced their OTR by -11.3 relative to the OSE stocks not likely to hit the limit. While the coefficient is not significant, it is consistent with the descriptive results that the OTR fell for the high OTR stocks.

However, what is of most interest is the coefficient on the interaction terms for the market quality measures. While the sign of the estimates indicate a worsening of liquidity for high OTR stocks relative to low OTR stocks (higher spreads and volatility, lower depth), the magnitudes of the effects are small, and none of the coefficients are significant. Hence, we conclude that the fee introduction did not have a significant effect on market quality at the OSE.

4.3. Did market quality change outside the OSE?

We next investigate another potential effect of the OTR fee introduction. Namely, that traders shift their trading away from the OSE. If so, we would expect that liquidity improves on the other exchanges relative to the OSE. Note that this is an auxiliary hypothesis.

We look for relative changes in liquidity by adopting a similar difference in differences technique to the one in the previous section, using the same pre and post periods. We use the technique to compare trading at the OSE and trading outside of OSE. It is important to note that this specification does not rely on trading outside of the OSE being a (untreated) control. While the other exchanges did not introduce an OTR fee, when the OSE (the treatment sample)

²⁰The standard errors are calculated using the Arellano (1987) adjustment of the White (1980) type of standard errors. The estimation is done using the R library `p1m`. Calculation of standard errors is described in Croissant and Millo (2008).

Table 6 Estimates of Difference in Differences investigation of “Low OTR” vs “High OTR” stocks

Panel A: One covariate: firm size

	OTR	Quoted (Rel) Spread	Effective Spread	Realized Spread	Roll	RV	Depth
β_2 d(Post Period)	-13.489 (27.870)	40.229*** (0.102)	11.038*** (0.140)	40.229*** (0.102)	17.393*** (0.143)	9.203*** (0.160)	-2,410.343*** (89.908)
β_1 d(high OTR)	-3.574 (7.829)	6.265*** (0.152)	1.020*** (0.044)	6.265*** (0.152)	0.508*** (0.033)	-0.054 (0.052)	-95.279 (146.464)
δ Interaction	-11.331 (8.438)	0.014 (0.210)	0.055 (0.047)	0.014 (0.210)	0.008 (0.061)	0.016 (0.063)	-232.846 (180.377)
ln(Firm Size)	1.588 (1.568)	-2.059*** (0.001)	-0.533*** (0.007)	-2.059*** (0.001)	-0.790*** (0.006)	-0.402*** (0.009)	145.464*** (4.322)
Observations	5,710	7,061	4,977	7,061	2,703	4,746	6,746
Adj. R ²	0.554	0.745	0.717	0.745	0.487	0.390	0.412

Panel B: Three covariates: firm size, RV, and 1/Price

	OTR	Quoted (Rel) Spread	Effective Spread	Realized Spread	Roll	Depth
β_2 d(Post Period)	32.652* (17.789)	37.532*** (0.662)	9.959*** (0.246)	37.532*** (0.662)	13.540*** (0.433)	-2,632.249*** (364.338)
β_1 d(high OTR)	10.065*** (2.914)	6.192*** (0.114)	0.965*** (0.042)	6.192*** (0.114)	0.438*** (0.029)	-284.537*** (95.276)
δ Interaction	-10.909 (7.797)	0.026 (0.116)	0.024 (0.040)	0.026 (0.116)	0.037 (0.051)	-47.826 (50.943)
ln(Firm Size)	-0.953 (0.950)	-1.939*** (0.032)	-0.484*** (0.012)	-1.939*** (0.032)	-0.618*** (0.019)	152.459*** (20.280)
RV	-1.716** (0.709)	0.196*** (0.059)	0.086*** (0.024)	0.196*** (0.059)	0.320*** (0.035)	-0.142 (7.557)
1/Price	-0.850* (0.451)	0.323** (0.133)	0.202*** (0.036)	0.323** (0.133)	0.301*** (0.085)	62.509*** (18.849)
Observations	4,715	4,715	4,664	4,715	2,700	4,619
Adj. R ² .	0.572	0.798	0.758	0.798	0.565	0.476

The coefficients are estimates of the regression $y = \beta_0 + \beta_2 d_{time} + \beta_1 d_{treated} + \delta d_{treated} \times d_{time} + \alpha X + \varepsilon$, where y is the various liquidity measures. $d_{treated}$ (d (high OTR)) is a dummy variable for treatment, where treatment is proxied by the maximal OTR in January-April 2012 being above 50. d_{time} (d (Post Period)) is equal to one if the observations is in the second period (September 2012) and zero otherwise. X are additional covariates. Each column specifies a regression. The dependent variables are shown in the column titles. The analysis is performed for the the order-to-trade ratio, quoted (relative) spread, the effective spread, the realized spread, the Roll measure, the realized volatility, and the depth. The spread measures, the Roll measure and the realized volatility are in percent. Depth is in thousands. As additional covariates we use the natural log of the market value of the firm’s equity (Panels A and B), and realized volatility and the inverse of the stock price (Panel B). The regressions include time and stock fixed effects (not reported). The standard errors are adjusted for clustering, using the Arellano (1987) adjustment of the White (1980) type of standard errors, as described in Croissant and Millo (2008). Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

did, it may still be the case that the control sample (trading outside of OSE) was affected indirectly. The difference in differences should rather be interpreted as looking at the relative change between liquidity at the OSE and liquidity of the same stocks at the alternative exchanges.

The results of this analysis are shown in Table 7. Again, we focus on the interaction term. Two of the estimates are significant at the 10% level, showing that the quoted spread and relative volatility both fell at the OSE relative to the alternative exchanges. With the exception of the realized spread, the sign of the other liquidity measures are also consistent with an improvement in liquidity at the OSE relative to the other exchanges, but these coefficients are not significant. Overall, these results provide additional support for our main result that the liquidity at the OSE did not deteriorate after the introduction of the fee, also when compared to the liquidity in OSE stocks at other exchanges.

Table 7 Difference in Difference OSE vs. largest alternative exchange

	Quoted (Rel) Spread	Effective Spread	Realized Spread	Roll	Realized Volatility	Depth
β_2 d(Post Period)	0.056 (0.045)	0.004 (0.023)	-0.009 (0.014)	0.031 (0.022)	0.021 (0.041)	-1.463 (20.494)
β_1 d(OSE)	-0.158*** (0.041)	-0.034** (0.016)	-0.028*** (0.010)	0.059** (0.024)	0.234*** (0.061)	156.978*** (32.247)
δ Interaction	-0.091* (0.048)	-0.016 (0.016)	0.018 (0.012)	-0.047 (0.034)	-0.125* (0.074)	32.677 (32.422)
Observations	2,758	2,732	2,798	2,206	2,674	2,760
Adj. R ²	0.593	0.467	0.128	0.077	0.046	0.723

The coefficients are estimates of the regression $y = \beta_0 + \beta_2 d_{time} + \beta_1 d_{treated} + \delta d_{treated} \times d_{time} + \varepsilon$, where y is the various liquidity measures. $d_{treated}$ ($d(OSE)$) is a dummy variable for treatment, where treatment is proxied by trading at the OSE, and non-treated is trading outside of the OSE. d_{time} ($d(Post Period)$) is equal to one if the observations is in the second period (September 2012) and zero otherwise. Each column specifies a regression. The dependent variables are shown in the column titles. The analysis is performed for the quoted (relative) spread, the effective spread, the realized spread, the Roll measure, the realized volatility, and the depth. The spread measures, the Roll measure and the realized volatility is in percent. Depth is in thousands. The regressions include time and stock fixed effects (not reported). The standard errors are adjusted for clustering, using the Arellano (1987) adjustment of the White (1980) type of standard errors, as described in Croissant and Millo (2008). Significance levels: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

5. Complementary analyses

The results in the previous section show little evidence of any negative effects on liquidity at the OSE related to the fee introduction. In this section, we complement these results by some additional analyses.

Firstly, we look at the monthly OTR threshold of 70:1 set by the OSE. Was this a threshold that would “bite” for the traders on the exchange? Secondly, we compare OTRs for individual traders pre- and post-introduction. Thirdly, we investigate the potential movement of trading from OSE to the competing exchanges using the information share of the OSE.

5.1. Was the limit of 70:1 likely to be binding?

One potential explanation of the non-result in the previous section is that the OTR constraint was set too high, such that it was not likely to be binding. If so, the introduction of the OTR was a “non-event” and we would not expect to see much reaction from traders. To investigate this we look at the OTRs in the pre-period (May 2012 - September 2012). We show examples of OTRs way above 70:1, both at the stock level and for individual traders.


5.1.1. Whole market

We first calculate the OTR for the whole market by, for each stock, counting the number of messages (order submissions, order withdrawals, and order modifications) to the exchange’s limit order book and dividing by the total

number of executed trades in the stock. Note that this is measured across all traders in each stock. This number represents a lower bound for the OTR for the more active traders on the exchange.

As an example, we show the OTRs for Statoil, the largest and most actively traded company on the OSE. Figure 3 shows daily estimates of OTRs for Statoil for the two years (2010–2011) leading up to the introduction of the fee. Here we see that there are days when the market-wide OTR for Statoil is above 70:1. Given some variation in the trading strategies used by the traders at the OSE, some of the traders in Statoil on these days must have had an OTR significantly above 70:1.

Figure 3 Daily order-to-trade ratios for Statoil: 2010-2012:5



The figure show time series of order-to-trade ratios for the stock Statoil for the period 2010–2012:5. For each day, we sum the number of messages into the exchange's system for this stock, and divide it by the number of actual trades.

5.1.2. Individual traders

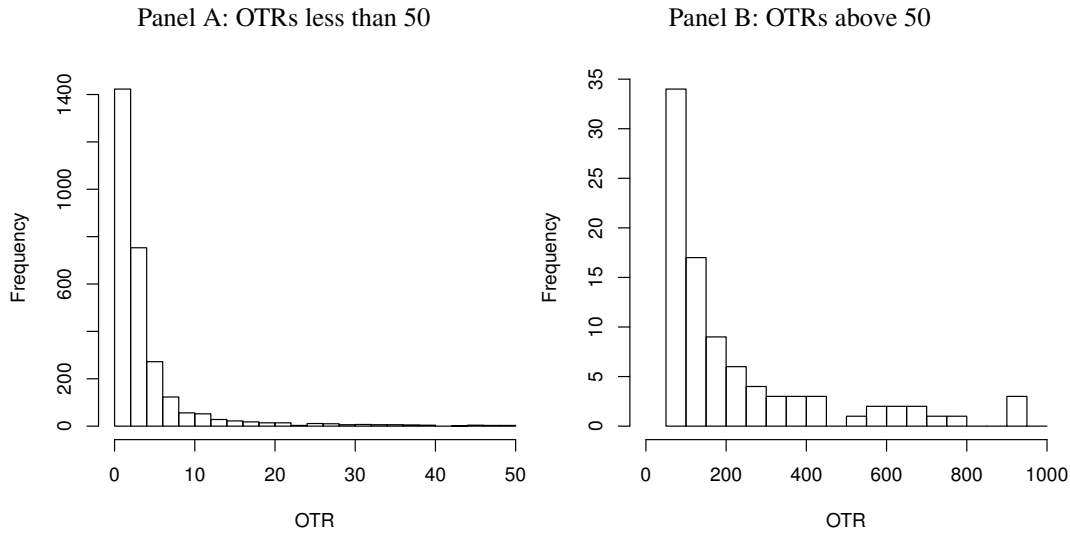
Since we can link each (anonymous) OSE member ID to trades and orders (messages into the limit order book), we can also calculate a monthly member-specific OTR by summing all trades and orders for that member during a month. To look at the situation before the exchange announces its new policy, we use data for May 2012. In Figure 4, we show the distribution of monthly OTRs calculated for all members and all stocks. The figure shows that the vast majority of OSE members have very low OTRs. However, we are mainly interested in the presence of high OTRs and therefore show (in Panel B) the distribution separately for OTRs above 50. As the figure shows, there is a nontrivial number of cases where OSE members have OTRs above 70:1. This includes some extreme examples. The highest monthly OTR is 944.

This shows that there were OSE members that would be likely to pay the fee if they did not adjust their behavior after the introduction of the OTR threshold.

5.2. Did individual traders react?

In this subsection, we examine whether OSE traders with high OTRs behave differently from other traders going forward. First, we consider traders that in May 2012 had high OTRs. If a trader makes no changes to his algorithms he runs the risk of paying the fee. If he instead modifies his algorithms to make paying the fee less likely we should

Figure 4 Distribution of OTRs for individual exchange members: May 2012



The figures show distributions of monthly OTRs for individual traders for May 2012. Panel A (on the left) show OTRs below 50. Panel B (on the right) show OTRs above 50. The OTR is calculated for each exchange member and each stock as the number of messages into the exchange's systems during a month divided by the number of actual trades in that stock by the member in that month.

see that this trader has a lower OTR in September 2012. In Panel A of Table 8, we show this in a simplified way. We calculate the percentage change in OTR from May 2012 to September 2012 separately for cases where a trader had an OTR above (below) 50 in May 2012. We find that in the high OTR cases, the average OTR was reduced by about 36% in September 2012, the first month the OTR was in place. In contrast, for the low OTR cases, the OTR increased with approximately 54%. This difference provides additional evidence that individual traders did react to the introduction of the fee.

Another way to make this point is to look at what happens in September 2012. Presumably, traders are watching the evolution of their aggregate OTR as the month progresses. If a trader sees that he currently has a high OTR, and hence may end up paying the fee, he will cut back on his order submission intensity for the remainder of the month. We investigate this prediction, again in a simplified way. We split the trades in September 2012 into two halves and calculate an aggregate OTR for each half. We then do the same calculations of differences as we did when comparing May 2012 to September 2012.

The idea is to look for evidence that high OTR traders lower the OTR for the second half of the month if they have high OTRs the first half. The results in Panel B of Table 8 confirm that in cases where the OTR was high in the first half of the month, traders react by a lowering their OTR in the second half of the month.

5.3. Information share analysis

Finally, we investigate whether there are major changes in price discovery relations between the OSE and its competing exchanges. The information share (IS), introduced by Hasbrouck (1995), attempts to identify which market places are important for information production or price discovery. To estimate the IS, Hasbrouck argues that the random walk component of a security price is the implicit efficient price, and defines the IS of a market as an attribution of the source of variation in the random-walk component to the innovations in that particular market. Unfortunately, in this estimation there is an extra degree of freedom, because the ordering of markets is arbitrary. Hasbrouck shows one can only estimate an upper and lower bound on the IS. We also estimate a competing measure to the IS, the

Table 8 Change in OTRs for individual traders

Panel A: May 2012 to September 2012

	Mean	Median	<i>n</i>
All stocks	51.1	0.0	2375
Stocks with OTR in May below 50	53.5	0.0	2311
Stocks with OTR in May above 50	-35.8	-50.7	64

Panel B: First to second half of September 2012

	Mean	Median	<i>n</i>
All stocks	27.1	0.0	7958
Stocks with OTR in first half of September < 50	27.9	0.0	7760
Stocks with OTR in first half of September > 50	-3.1	0.0	198

The tables present the percentage change in OTR for individual traders. Panel A uses changes from May 2012 to September 2012. For each trader and each stock, we calculate the (monthly) OTR as the total number of messages into the exchange's system during the month, divided by the number of trades. This calculation is done for May 2012 and September 2012, before calculating the percentage change in the OTR from May to September. We report the mean and median. We also calculate half-monthly OTRs using the same method. The calculation is done for the first and second half of September 2012, before calculating the percentage change in the OTR from the first to the second half of the month. Panel B shows the results. We present results for all stocks, stocks with an OTR below 50, and stocks with an OTR above 50.

common factor component (CS).²¹ The relative merits of the two measures are still disputed,²² but both measures are informative about the relative share of price discovery of two market places. We therefore also include the CS.

Using the IS/CS allows us to investigate the hypothesis that after the OSE introduction of the fee, quotes at the OSE have become more stale. If so, the share of price discovery happening at OSE should fall after the introduction of the fee. We investigate this by estimating the IS/CS for the stocks with significant trading outside of the OSE. For each stock, we calculate the price discovery metrics on a daily basis.²³ As before, we compare estimates for May 2012 to the same measures in September 2012. The results are in Table 9.

Table 9 OSE Information Share results

	Average		<i>t</i> -stat	<i>p</i> -value
	May	September		
IS (original ordering)	0.763	0.756	-0.730	0.465
IS (reversed ordering)	0.536	0.542	0.512	0.609
CS	0.550	0.539	-1.185	0.237

The table reports summary statistics for information share in April and September of 2012. We use stocks in the OBX index. For each stock, we calculate the IS and CS of OSE vs. its largest competitor exchange, which may be Stockholm, Chi-X, Turquoise, or London. We use five-second interval observations of the last transaction price. The time matching is based on the time when the transaction was reported in London (to Reuters), not the original time on each exchange. The IS (original ordering) is using OSE as the first exchange, and the IS (reversed ordering) has the opposite ordering. We list averages of the IS in May and September. We also provide the results of a *t*-test of equality of means across the two months. Calculation uses the `pshare` routine in the `ifrogs` R library.

The interpretation of each of these measures is the fraction of price discovery due to a given market. So the IS of 0.763 in Table 9 should be interpreted as 76.3% of the price discovery happening at the OSE. The estimates show that the majority of price discovery of OSE listed stocks is at the OSE both before and after the introduction of the fee. Moreover, none of the estimates change significantly from May 2012 to September 2012, which provides further support that the market quality at the OSE was *not* negatively affected by the fee introduction.

²¹The component share (CS) was introduced by Harris, McInish, Shoesmith, and Wood (1995), Booth, So, and Tse (1999), and Chu, liang Gideon Hsieh, and Tse (1999), as an application of results in Gonzalo and Granger (1995).

²²See the special issue of the *Journal of Financial Markets* (Lehmann, 2002), and the more recent paper by Yan and Zivot (2010).

²³The IS/CS is calculated using five-second interval observations of the last transaction price. The time matching is based on the time when the transaction was reported in London (to Reuters), not the original time on each exchange.

6. Conclusion

In this paper we examine the effect of introducing a fee specifically targeted at some HFT strategies. The fee is payable by traders whose messages to the exchange's limit order book exceed 70 times the number of actual trades these orders resulted in. Such a fee has been proposed by several policy makers and market participants as a way for the most active users of the exchange's infrastructure to pay some of the costs they are imposing on the system operator and other traders.

The fee was introduced at the OSE on September 1, 2012. The risk of introducing a fee based on an OTR is that it may negatively affect the quality of trading at the exchange. We investigate the issue in two ways. First, we examine the evolution of a number of measures for market quality such as depth, spread, trading costs, volatility, and turnover. Over the longer period of a year around the introduction of the fee, we find that market quality actually improves at the OSE relative to its competitors. However, when we examine the introduction of the OTR fee, through various difference in differences specifications, we also find that there is no evidence of a change directly linked to the introduction of the fee.

The same conclusion is reached through our second type of analysis, a comparison of the price discovery of the OSE relative to its competitors. We find that the share of OSE in price discovery, using both the information share and common factor component measures, remains stable throughout the period.

The experience of the OSE seems to be very different from the experience when a similar measure was introduced in Italian and Canadian equity markets. In both of these cases, we saw a decline in liquidity around their introduction of a similar type of fee. At the OSE, we find no effect. This may be due to differences in the design of the fees, where the OSE design is more aimed at encouraging liquidity provision, such as HFT market making strategies.

However, an alternative explanation may be that the OTR fee of the OSE was too lenient, and nobody needed to change their behavior. To show some evidence against this explanation, we examined all individual traders at the OSE and showed that some of them had OTRs before the introduction of the fee significantly above levels that would have triggered the fee. We also examined the same traders after the fee was introduced and show that they did react by lowering their OTRs. We also show that traders changed their behavior towards the end of a month to offset a high OTR at the beginning of the month. Both these results show that the OTR regulation was not trivial, and that traders needed to change their algorithms.

To conclude, our results show that the OSE seems to have achieved its goals, as the exchange achieved a reduction in message traffic from the most extreme HFTs without an exodus of traders and with the market quality intact. In fact, the exchange tells us that no trader on the exchange has yet paid the fee.

References

- James J Angel, Lawrence Harris, and Chester S Spatt. Equity trading in the 21st century. *Quarterly Journal of Finance*, 1(1):1–53, 2011.
- Manuel Arellano. Computing robust standard errors for within-group estimators. *Oxford Bulletin of Economics and Statistics*, 49(4):431–34, 1987.
- Robert Battalio, Shane A Corwin, and Robert Jennings. Can brokers have it all? On the relation between make-take fees and limit order execution quality. *Journal of Finance*, LXXI:2193–2238, 2016.
- Bruno Biais, Thierry Foucault, and Sophie Moinas. Equilibrium fast trading. *Journal of Financial Economics*, 116(2):292 – 313, 2015.
- Øyvind Bøhren and Bernt Arne Ødegaard. Patterns of corporate ownership: Insights from a unique data set. *Nordic Journal of Political Economy*, 27:57–88, 2001.
- G Geoffrey Booth, Raymond W So, and Yiuman Tse. Price discovery in the German equity index derivatives markets. *Journal of Futures Markets*, 19(6):619–643, 1999.
- Jonathan Brogaard, Bjorn Hagströmer, Lars Nordén, and Ryan Riordan. Trading fast and slow: Colocation and liquidity. *Review of Financial Studies*, 28(12):3407–3427, December 2015.
- Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4):1547–1621, 2015.
- Bidisha Chakrabarty, Pankaj K Jain, Andriy Shkilko, and Konstantin Sokolov. Speed of market access and market quality: Evidence from the SEC naked access ban. Working paper, SSRN, May 2014.
- Quentin C Chu, Wen liang Gideon Hsieh, and Yiuman Tse. Price discovery on the S&P 500 index markets: An analysis of spot index, index futures, and SPDRs. *International Review of Financial Analysis*, 8(1):21–34, 1999.
- Yves Croissant and Giovanni Millo. Panel data econometrics in R: the plm package, 2008. R vignette, available at CRAN.
- Thierry Foucault, Johan Hombert, and Ioanid Roşu. News trading and speed. *Journal of Finance*, 71(1):335–382, February 2016.
- Sylvain Friederich and Richard Payne. Order-to-trade ratios and market liquidity. *Journal of Banking and Finance*, 50:214–223, 2015.
- Jesus Gonzalo and Clive W J Granger. Estimation of common long-memory components in cointegrated systems. *Journal of Business and Economic Statistics*, 13(1):27–35, 1995.
- Björn Hagströmer and Lars Nordén. The diversity of high frequency traders. *Journal of Financial Markets*, 16: 741–770, 2013.
- Frank H deB Harris, Thomas H McNish, G Shoesmith, and Robert A Wood. Cointegration, error correction and price discovery on informationally-linked security markets. *Journal of Financial and Quantitative Analysis*, 30:563–579, 1995.
- Joel Hasbrouck. One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4):1175–1199, 1995.
- Terrence Hendershott, Charles M Jones, and Albert J Menkveld. Does algorithmic trading improve liquidity? *Journal of Finance*, 66(1):1–33, 2011.
- Gregory Laughlin, Anthony Aguirre, and Joseph Grundfest. Information transmission between financial markets in Chicago and New York. *Financial Review*, 49:283–312, 2014.

- Bruce N Lehmann. Some desiderata for the measurement of price discovery across markets. *Journal of Financial Markets*, 5(3):259 – 276, 2002.
- Michael Lewis. *Flash Boys*. Allen Lane, 2014.
- Katya Malinova and Andreas Park. Subsidizing liquidity: The impact of make/take fees on market quality. *Journal of Finance*, LXX(2):509–537, April 2015.
- Katya Malinova, Andreas Park, and Ryan Riordan. Taxing high frequency market making: Who pays the bill? Working paper, University of Toronto, June 2016.
- Tom Grimstvedt Meling. Anonymous trading in equities. Working Paper, University of Bergen, June 2016.
- Tom Grimstvedt Meling and Bernt Arne Ødegaard. Tick size wars, high frequency trading, and market quality. Working Paper, University of Stavanger, June 2017.
- Albert J Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16:712–740, 2013.
- Albert J Menkveld. The economics of high-frequency trading: Taking stock. In Andy Lo and Robert Merton, editors, *Annual Review of Financial Economics*, volume 8 of *Annual Review of Financial Economics*. Annual Reviews, 2016.
- Albert J Menkveld and Marius A. Zoican. Need for speed? Exchange latency and liquidity. *Review of Financial Studies*, 16:712–740, 2016.
- Randi Næs and Johannes A Skjeltorp. Order book characteristics and the volume–volatility relation: Empirical evidence from a limit order market. *Journal of Financial Markets*, 9:408–432, 2006.
- Randi Næs, Johannes A Skjeltorp, and Bernt Arne Ødegaard. Stock market liquidity and the Business Cycle. *Journal of Finance*, LXVI:139–176, February 2011.
- U.S. Securities and Exchange Commission. Report to congress on decimalization. SEC Report, July 2012.
- U.S. Securities and Exchange Commission, Division of Trading and Markets. Maker-taker fees on equities exchanges. Memo to Equity Market Structure Advisory Committee, October 2015.
- U.S. Securities and Exchange Commission, Division of Trading and Markets. Certain issues affecting customers in the current equity market structure. Memo to Equity Market Structure Advisory Committee, January 2016.
- Halbert White. A heteroskedasticity-consistent covariance matrix estimator and a direct test of heteroskedasticity. *Econometrica*, 48:817–38, 1980.
- Bingcheng Yan and Eric Zivot. A structural analysis of price discovery measures. *Journal of Financial Markets*, 13(10):1–19, 2010.