

Equity trading costs have fallen less than commonly thought. Evidence using alternative trading cost estimators

Valeriia Klova and Bernt Arne Ødegaard*

Mar 2019

Abstract

Equity markets are evolving rapidly. The technology of financial intermediation has changed from monopolistic (manual) market makers to multiple algorithms providing liquidity in competing order books, both visible and dark. Has this changed market quality? To answer this, we need measures of market quality (liquidity) invariant to technological innovation in intermediation. In particular, innovation has led to a huge drop in order sizes due to order splitting. Orders are spread out both in time and across exchanges. We use data from the US and Norway to show that the last two decades' marked fall in average spreads is driven by the decline in transaction sizes. Using alternative estimators of transaction costs less sensitive to trade size, such as the Corwin and Schultz (2012) and Abdi and Rinaldo (2017) high/low estimators, we show that equity market quality has improved less than commonly thought.

Keywords: Equity Trading Costs, Spread, High/Low Estimator

JEL Codes: G10; G12; G20; G23

The world's trading systems for equities are evolving rapidly. Yesterdays monopolistic market maker has been replaced with competing algorithms providing continuous liquidity in numerous separate order books (both visible and dark). Is this system an improvement? From a policy perspective, the key question is whether the changes to markets have benefitted the end users: investors trading equities, be they individuals, or mutual funds trading on behalf of the same individuals. To investigate this it is necessary to look at time series of measures of market quality, such as trading costs, for these users.

The market microstructure literature has provided numerous empirical liquidity measures.¹ Liquidity is a concept with different connotations, where quantity, costs and time are different aspects of liquidity. Which of these measures are most relevant for the end users? What are desirable *properties* of such measures? We argue that a key property is invariance to changes in the technology of financial intermediation. Several common measures of liquidity may be affected by changes in market

*Klova (valeriia.klova@uis.no) is a PhD candidate at the University of Stavanger (UiS) Business School. Ødegaard (bernt.a.odegaard@uis.no) is a professor at the University of Stavanger (UiS) Business School and an adjunct professor at the Norwegian School of Economics (NHH). We are grateful for comments from Sean Foley, Kenneth Kavajecz, Tom Meling, Randi Næs, and seminar participants at the FIBE conference in Bergen and at the University of Stavanger. We are grateful to Thomas Borchrevink and Christian Ringstad at the Oslo Stock Exchange for access to OSE trade data. We thank Tom Meling for help with data on lot size changes at the OSE.

¹See Holden, Jacobsen, and Subrahmanyam (2013) for a recent survey.

structures driven by innovation in intermediation. These innovations may not be particularly relevant for the end user. Let us discuss a couple of examples.

First, consider the quantity dimension. There are several aspects of the new financial market place that inflate quantities. One is that inventory management has moved onto the order book. The monopolistic market maker would allow for sizeable inventory imbalances during the trading day. The modern algorithmic market makers are more active in unloading inventory immediately.² Another concerns the interactions between market places. The process of aligning prices across exchanges may induce arbitrage trading of the same stock at several exchanges. These are both examples of trading between financial intermediaries, not end users. We have over time seen that a given quantity of trading by end users is supported by increased trading between financial intermediaries. An increase in pure quantity measures, such as stock turnover, is thus not necessarily an improvement from the point of view of the end user. These issues also affect other measures that include quantities, such as the Amihud (2002) measure.

Second, consider the cost dimension. Another aspect of the new market environment is that traders to a much larger degree than before rely on splitting orders, both across exchanges and across time, in an attempt to minimize price impact. These changes to trading strategies have led to a dramatic decrease in average trade sizes.³ This has consequences for spread types of measure. When one measures a spread, it is measured using either an actual transaction, or the best bid and ask in the limit order book. In both alternatives there is a *quantity* involved, the number of shares. Hence, a bid-ask spread estimate is really an estimate of an average transaction cost *for trading an average number of shares*.

Hence, changes in both volume-based measures, such as the well known Amihud measure, and spread measures, be they quoted, effective or realized spreads, may be driven by how the advent of High Frequency Trading (HFT) has affected the trading process. Maureen O'Hara (2015) argues that "the new high frequency world" should lead to a rethink of research in market microstructure. That is the context of this paper. Which empirical measures of market quality are most relevant for measuring the evolution of market quality?

We point to an alternative class of estimators, recently developed, that relies on the difference between the highest and lowest trade price during a trading day. This class of estimators relies on a simple intuition: the highest price during the day is most likely to have been the result of a price improving trade, where the buyer was the active part in pushing the price (and crossing the spread). The highest price is therefore likely to have been at the ask. The opposite argument gives that the lowest price is likely to have been at the bid. The difference between the highest and lowest price therefore includes one (implicit) spread. The Corwin and Schultz (2012) and Abdi and Rinaldo (2017) proxies are alternative ways of estimating this spread. While these estimators were developed mainly to obviate the necessity of microstructure data, we argue that this class of estimators also has the desirable property that it is less sensitive to changes in market structure.

Let us first show an example, where we compare the time series evolution of these high-low measures with a spread measure and the Amihud illiquidity measure. Figure 1 shows two example

²Menkveld (2013) shows that high-frequency market makers hold their inventories for a very short period of time.

³See e.g (Angel, Harris, and Spatt, 2011, pg 16) and (O'Hara, 2015, pg 260).

companies: General Electric, traded at the New York Stock Exchange, and Norsk Hydro, traded at the Oslo Stock Exchange, over the last thirty years. These example companies are among the largest (and most liquid) stocks on their respective exchanges. The spread measure and the Amihud measure fall markedly during the period, while the other two measures do not have the same marked fall. The figure makes clear the importance of our research for policy purposes. Conclusions about market quality are very different depending on which liquidity measure is deemed most relevant.

In this paper we concentrate on one aspect of the new trading environment, the reduction in transaction size. From the point of view of the end user, to evaluate market structure, a cost/liquidity measure needs to be invariant to the average order size. We therefore investigate the sensitivity of alternative liquidity measures to order size.

In our analyses we use data for two different markets. We first use data for the NYSE. We show time series of various trading cost estimators and compare their evolution. We demonstrate that the behaviour we illustrated for GE is true for the whole exchange, spreads fall, while other trading cost measures are essentially unchanged. Unfortunately, the US markets are not best suited to studying the effect of order size. The main issue with the US market place(s) is that trading has fragmented substantially during the period, where the main listing markets (NYSE, NASDAQ) have been complemented with numerous alternative markets. The different market places cater to different types of customers, with different trading needs. So one would need to look across many markets to evaluate the typical order size.

In contrast, trading in Norwegian shares was concentrated at the Oslo Stock Exchange for much longer, fragmentation chiefly starting with the introduction of MiFID in 2008. The OSE is still the main market place for the OSE-listed companies. The OSE therefore offers a long period with trading where we can identify order sizes for the complete market.

We start our analysis by pointing out and establishing carefully the degree of dichotomy between the various transaction cost estimates. This is done using data for both the NYSE and the OSE.


Next, in an attempt to understand what is driving the difference in estimates, we investigate whether the order size affects the various estimates to the same extent using daily transaction data for the OSE over the period 1999–2016. We hypothesize that order size is positively related to the (closing) relative and effective spreads and unrelated to the high-low estimators and the Amihud measure. Further, we hypothesize that the gap between spreads and high-low estimates can be explained by order size. The order size hypothesis is strongly supported by our results. Lower order sizes leads to lower closing and effective spreads, while they do not have such effect on the high-low estimators and the Amihud measure. Moreover, the difference between spreads and high-low estimators is to a large extent driven by the change in order size. The difference between these estimators increases significantly as the order size decreases.

Our first analysis is essentially a joint time series estimation, which is potentially vulnerable to joint trends. We therefore perform several identification exercises leveraging two types of exogenous events. The first looks at the times when a stock splits. After a stock split, if the (dollar/krone) transaction size is unchanged, the number of shares traded increases. We therefore use stock splits as a source of exogenous variation in transaction sizes. Another interesting source of exogenous


Figure 1: Liquidity estimates for General Electric (GE) and Norsk Hydro

The figure shows quarterly estimates of liquidity for General Electric (GE) (Panel A) for the period 1995 to 2016, and for Norsk Hydro (Panel B) for the period 1980–2017. Spreads (toward the left axis) are estimated using three methods: *AR* is the Abdi and Ranaldo (2017) spread estimate, *CS* is the Corwin and Schultz (2012) spread estimate, *Spread* is the relative spread, measured as the difference between closing bid and ask divided by the mid-price. The spread measures are calculated on a daily basis and averaged across days in a quarter. Toward the right axis we plot quarterly estimates of the Amihud ILLIQ measure.

Panel A: GE



Panel B: Norsk Hydro



variation is the OSE's use of a minimum lot size. In the early part of our period the OSE maintained a list of minimum lot sizes. This lot size was revised on a stock-by-stock basis a few times a year by the exchange. We use times of changes in lot size as exogenous variation in transaction size in a similar manner to stock splits.

In both cases we look at the changes to the various liquidity measures around the events, and set up difference-in-differences analyses to directly estimate the presence of a causal effect, and its magnitude. The analyses of stock splits and lot size changes show that the high-low estimators and the Amihud measure are not affected by the changes in average transaction size and minimum order size, while the closing and effective spreads are significantly affected. This finding corroborates the story of a less pronounced effect of changes in order size on the high-low estimators compared to the conventional spreads.

Our research is not alone in considering how to best measure liquidity in the new market environment. Angel et al. (2011), Chordia, Roll, and Subrahmanyam (2011) and Holden and Jacobsen (2014) all question how our existing liquidity measures fare in the new fast trading environment. The search is on for new liquidity measures. An interesting example is the measure of Barardehi, Bernhardt, and Davies (2018). They propose to modify the Amihud measure by calculating it over varying time intervals of fixed-dollar volumes, which is an interesting way of accounting for changes in trading intensity.

Our research has the same starting point. We do not propose a modified liquidity measure, we instead point to recently introduced high-low measures and investigate their properties. We argue that the high-low estimators should be regarded as superior measures of transaction costs in modern markets. The supporting evidence also comes from the results for the Amihud measure, which is arguably an accurate measure of price impact regardless of order size. Both the high-low estimators and the Amihud measure are shown to be unaffected by order size, which makes them robust measures of transaction costs. However, the Amihud measure may have issues with inflated aggregate quantities due to increased trading between intermediaries. The high-low estimators should be preferred as they capture transaction costs to a larger degree, i.e. they also include the spread component in addition to price impact.

Let us finally point out the important implication of our results: transaction costs, as measured by the high-low estimators, have not fallen, which contradicts the common belief that liquidity has improved in the new HFT world.

The paper proceeds as follows. Section 1 gives an overview of various liquidity/trading cost estimators and discusses the effects of order size. Section 2 describes the data and compares time-series development of liquidity estimators. Section 3 provides results from a regression analysis of how order size affects trading costs. Section 4 analyses exogenous variation in transaction size by investigating stock splits and lot size changes. Section 5 concludes.

1 Measuring equity trading costs

In this study we are concerned with the cost of trading equities in stock markets. Using a classification from Keim and Madhavan (1998), total trading costs can be split into explicit components such as

broker’s fees, and implicit components such as spread costs, price impact costs, and costs related to delayed or uncompleted trading (opportunity costs). The magnitude of these components are found to vary with factors such as trade difficulty, order submission strategy, market design and investment style. Hence, estimation of total trading costs for a specific investor is hard, and requires detailed data on the entire order submission process.

In this paper we abstract from the identified investor perspective. We instead look at estimates of equity trading costs that can be calculated from publicly observable data. The magnitude of such measures can be said to be informative about the expected transaction costs faced by the *typical investor*, i.e. a retail investor buying a small number of shares. It is less clear how informative it will be for institutional investors who want to transact larger quantities.

The empirical finance literature has suggested numerous transaction cost estimates. We will first define them before discussing their relative merits.

1.1 Trading cost estimators

The benchmark measure of transaction costs is the effective bid-ask spread, which is defined as the difference between the actual transaction price and the mid-price. The formula for the effective spread (in relative terms) $s_{eff,jt}$ for a stock j for transaction t is

$$s_{eff,jt} = q_{jt} * \frac{p_{jt} - m_{jt}}{m_{jt}}, \quad (1)$$

Here p_{jt} is a transaction price, m_{jt} is a mid-price, and q_{jt} is an indicator variable that denotes direction of trade,

$$q_{jt} = \begin{cases} +1 & \text{(buyer-initiated)} \\ -1 & \text{(seller-initiated)} \end{cases}. \quad (2)$$

Calculation of the effective spread requires transaction- and order-level data. In the absence of signed transaction data, the classification of transactions is typically performed using the Lee-Ready algorithm or other classifying methods, which are not free of errors and, subsequently, may introduce bias in the effective spread (Foucault, Pagano, and Röell, 2013). The effective spread is an acknowledged benchmark for measuring transaction costs in the microstructure literature. It is common to compare performance of other measures against the effective spread.

Another widely-used measure of transaction costs is the closing or quoted relative bid-ask spread. The closing spread is a straightforward measure compared to the effective spread. It is simply defined as the difference between the best ask and bid (closing) quotes, divided by the mid-price:

$$s_{closing,jt} = \frac{ask_{jt} - bid_{jt}}{m_{jt}} \quad (3)$$

The closing spread is an accurate measure of implicit transaction costs for small transactions that can be filled at the best quotes.

Supplementing these spread measures, a number of alternative measures relying on low-frequency data has been developed. The first such measure was Roll (1984). The intuition of the Roll measure

is that the existence of a spread will have implications for the time series behaviour of prices. As transactions prices bounce between the bid and ask, one can use the autocorrelation between subsequent transaction prices to infer the implied size of the bid/ask spread. While the Roll measure was developed for transaction prices, it is also often applied to daily (closing) prices.

Another transaction cost measure was developed by Lesmond, Ogden, and Trzcinka (1999) (LOT). This measure uses dates without transactions to infer the magnitude of transaction costs that inhibits price movements. The idea is that the market movement on a given day predicts a price change in a stock. If there is no transaction in the stock moving the stock price, the implicit transaction costs must be larger than the predicted stock price change. Lesmond et al. develop a Maximum Likelihood estimator of the implicit transaction costs. A problem with the LOT measure, and other measures that rely on zero returns (Fong, Holden, and Trzcinka, 2017) is that the likelihood of observing an exact zero return has declined as tick sizes and order sizes have fallen (Mazza, 2015).

As a significant part of transaction costs is price impact, a number of price impact measures have been developed. These measures capture price impact of trade, i.e. the change in price as a result of a transaction. The most common low-frequency measure of price impact is the Amihud (2002) measure. The Amihud measure greatly resembles the Kyle (1985) lambda measure of price impact. It is calculated as the ratio of absolute return to trading volume, averaged over a given time period. The Amihud measure is estimated on a daily basis and therefore captures day-long price impact. It could be further averaged to obtain the estimates of price impact at lower frequencies. The main issue with the Amihud measure is that it cannot be defined for days with no trading.

$$Amihud_t = \frac{1}{N} \sum_N \frac{|Ret_t|}{TradeVol_t} \quad (4)$$

More recently, other measures of transaction costs have been proposed that utilize another piece of information typically collected on a daily basis, the highest and lowest transaction prices during a day. The intuition that these measure rely on is that most likely the lowest price will be seller initiated, and therefore traded at a bid prices. Conversely, the highest price will most likely have been a result of crossing the spread from below. Therefore, the difference between the highest and lowest price will include one (implicit) spread. This intuition is used in different ways in the two estimators we consider.

The first measure, Corwin and Schultz (2012) (CS), additionally assumes that the variance component of the price process, the high-low ratio, is proportional to the time interval, which allows them to use a variance ratio of two-to-one day price variation to back out the implicit spread. The second estimator, Abdi and Rinaldo (2017) (AR), uses a different approach. Again, the estimator relies on the difference between the high and low price including one spread. But the intuition used in constructing the estimator is that the average of the high and low price, the mid-range, must have half the spread on either side, and be an unbiased estimate of the “true” price. The estimator then uses similar intuition as used by Roll (1984), that the bounce midpoint-close-midpoint can be used to infer the implied spread.

In calculating these estimates there are a number of implementation issues. First, three of the

above estimators (Roll, CS and AR) can result in negative values for the estimated spread. We follow the literature and set the spread to zero in these cases.⁴ Second, the Corwin and Schultz (2012) estimator is adjusted for overnight price movements, by potentially correcting for the difference between the close the first day and the open the next day.⁵

Despite the aforementioned issues with estimation of the high-low spread, Corwin and Schultz (2012) show that it is significantly better at capturing cross-sectional variation and month-to-month changes in effective spreads than alternative low-frequency spread estimators, namely, the Roll (1984) measure, the effective tick estimator of Holden (2009), and the LOT measure of Lesmond et al. (1999). It also outperforms other low-frequency measures in capturing time-series variation in individual stock spreads, with an exception for the largest stocks.

The performance of the high-low spread estimator of Corwin and Schultz (2012) has been investigated in a number of recent microstructure studies. To test the accuracy of the Corwin and Schultz (2012) high-low spread estimator, Lin (2014) performs simulation analyses under different scenarios and finds that its accuracy increases in spread size and transaction frequency, and decreases in stock volatility. This is consistent with the arguments outlined by Corwin and Schultz (2012) regarding a possible bias in high-low spreads for infrequently traded stocks and for stocks that exhibit high volatility over a 2-day period.

A more elaborate analysis of the high-low measure's accuracy is conducted by Bleaney and Li (2015). They question the results obtained in Corwin and Schultz (2012) under their nearly perfect conditions, and construct simulations that account for the numerous possible departures from the ideal case: time-varying spreads, autocorrelation in midpoint returns and trade directions, price impact, and feedback trading. They also examine the effect of sampling frequency on the accuracy of various spread estimates, including the high-low ones. For example, they report that the high-low estimator based on low-frequency data is consistent and reliable, and outperforms other low-frequency estimators. The high-low estimator however performs worse than the Roll (1984) estimator and the Huang and Stoll (1997) estimator when high-frequency data is used.

Karnaukh, Ranaldo, and Söderlind (2015) test the accuracy of various low-frequency spread estimators, including the high-low estimator, in foreign exchange markets. Among the spread estimators they utilize⁶, the high-low estimator provides the highest correlation with the effective spread benchmark. Schestag, Schuster, and Uhrig-Homburg (2016) additionally show that the high-low estimator performs as well as the Roll and Gibbs estimators in bond markets.

The AR high-low spread estimator is also shown to outperform alternative low-frequency spread estimators, including the CS estimator (Abdi and Ranaldo, 2017). Specifically, it exhibits the second highest time-series correlation (after the quoted spread) and the highest cross-sectional correlation with the effective spread benchmark. Abdi and Ranaldo additionally show that their estimator is only

⁴For discussion of the Roll estimator see Harris (1990). For the CS estimator, Corwin (2014) argue that the best solution is to set negative 2-day spreads to zero before calculating monthly averages. For the AR estimator, Abdi and Ranaldo (2017) also argues that imposing zero is the preferred method.

⁵The simulation results in Corwin and Schultz (2012) indicate that spreads adjusted for overnight returns are more accurate than non-adjusted spreads.

⁶The alternative low-frequency spread estimators investigated in Karnaukh et al. (2015) are the Roll (1984) estimator, the Gibbs estimator of Hasbrouck (2004), the effective tick estimator of Holden (2009), the LOT and the Zeros measures of Lesmond et al. (1999), and the FHT measure of Fong et al. (2017).

marginally sensitive to the number of trades per day and as a result provides a higher explanatory power over spreads for less liquid stocks compared to the CS estimator.

To summarize, the high-low spread estimators use daily high and low (and close) prices, while spreads use quotes and transactions. The microstructure literature suggests that both the Corwin and Schultz (2012) and Abdi and Rinaldo (2017) high-low spread estimators yield spread estimates comparable to the effective spread, as implied by the high degree of (cross-sectional) correlation between them. Furthermore, they tend to significantly outperform other low-frequency spread estimators.

1.2 On trading size

As shown by e.g. Angel et al. (2011), the average NYSE trade size has been steadily falling.⁷ However, since order splitting algorithms overtook the market, a small trade does not convey the same information as before. That is, a small order could be a part of a larger order split across time and trading venues. When this is the case, the closing and effective spreads may not be the best measures, as they can make you erroneously believe that transaction costs have been going down due to lower price impact associated with smaller trades. In other words, the benchmark spreads capture only the top of the iceberg, neglecting the possibility that a small order could be a part of a large split order and potentially have much bigger price impact on the market. Theoretically, in order to calculate the effective spread for a split order, the average price over the entire order should be obtained and compared to the initial mid-price. However, tracing split orders is practically impossible from tick type of data, which makes estimation of total price impact of split orders unfeasible.

The high-low spread estimators of Corwin and Schultz (2012) and Abdi and Rinaldo (2017) eliminate the aforementioned issues with the benchmark spreads and therefore are arguably better measures of transaction costs for split orders. First, the high-low estimators derive the spread using the daily high and low (and close) prices, and thus are not tied to any particular order size. This implies that they are less sensitive to order size changes, e.g. order splitting. Second, as pointed out by Corwin and Schultz and Easley, de Prado, and O'Hara (2016), the high-low spread estimators may be more comprehensive measures than the effective spread as they capture price impact not only due to large orders, but also due to succession of buy or sell orders through the high-low price ranges.

In this paper we therefore investigate whether the average order size has heterogeneous effects on alternative spread estimators. We hypothesize that the order size will have a positive effect on the benchmark spreads, but no significant effect on the high-low estimators. As previously argued, the high-low estimators are expected to be less sensitive to order size changes, and consequently, represent a more accurate way of measuring transaction costs for split orders. Similarly, the Amihud measure is hypothesized to be unaffected by order sizes as it is calculated using daily data.

1.3 Institutional trading costs

An alternative approach to investigate what has happened to transaction costs is to attempt to measure it with individual-level transaction data. A couple of recent papers, Eaton, Irvine, and Liu (2018)

⁷See figure 15 in Angel et al. (2011).

and Frazzini, Israel, and Moskowitz (2018), gives some insights into institutional trading costs in the new trading environment:

Their findings parallel ours. Eaton et al. (2018) use US transaction data from Abel/Noser to argue that the price impact of institutions is not correlated with common liquidity proxies, such as the bid/ask spread.

Frazzini et al. (2018) use trades of a single, large institutional investor across many markets. Among their many investigations they compare the time series evolution of their estimate of institutional trading costs with the effective spread. They show that while the effective spread is falling markedly in the period 2000-2015, the same is not happening for their estimate of transaction costs, which is not changing much (except for an increase around the financial crisis).⁸

2 Market places and data

In this section we discuss the data and market places.

2.1 Background on the Norwegian equity market

In our analysis we use detailed data on trading at the OSE. As this market is less known than the US markets, we start by giving some background for the market place.⁹ The OSE is a medium-sized stock exchange, and is ranked among the 30 largest equity markets in the world in terms of market capitalization. It is the only regulated securities market in Norway. Computerized limit order books were introduced to the OSE in 1982. The exchange functioned as a hybrid market, where orders could be crossed in the computerized limit order book, or manually. In 1999 the OSE became fully electronic, all trades are crossed in the exchange's computer systems. In the current regime, trading at the OSE is continuous during the day except at the open and close, where prices are cleared through an auction. The exchange operates on a price-time priority basis, as is common in limit order markets. The automation of trading was followed by improvement in computer infrastructure, and a gradual reduction in latency. For example, the OSE migrated to the Millennium trading platform in November 2012, which resulted in a decline in latency from 3 to 0.113 milliseconds.

For the early part of the period OSE was the only market place for most OSE listed securities. The only exceptions were a few of the largest Norwegian companies which had cross-listed in the US.¹⁰ The OSE lost its monopoly position following the introduction of the EU's Markets in Financial Instruments Directive (MiFID), which went into force on November 1, 2007. MiFID shares a common objective of inter-market competition with the Regulation National Market System (Reg NMS), adopted in US in 2006.

As a consequence of the new regulations, both the European and US equities markets have become highly fragmented. Proliferation of trading venues significantly expanded execution opportunities for

⁸Compare their figures 7 and 8.

⁹For further background, see e.g. Bøhren and Ødegaard (2001), Næs and Skjeltorp (2006), Næs, Skjeltorp, and Ødegaard (2008), Meling and Ødegaard (2017) and Jørgensen, Skjeltorp, and Ødegaard (2018).

¹⁰For example, the largest Norwegian company, Statoil (now Equinor), cross-listed at the NYSE. Telenor, the national telecom, cross-listed at NASDAQ.

large orders. Orders can be executed on “lit” venues (public limit order books) or in “Dark Pools” (less regulated market places). Investors can also employ order splitting algorithms to split large orders both across time and trading venues.

Despite aiming at the same objective, these two regulations have important differences in execution and order handling rules. One fundamental difference concerns the Reg NMS order protection rule against trade-throughs, which ensures execution of orders at the best available prices. In MiFID this is replaced with a more general best execution rule, which not only considers price, but also costs, speed, likelihood of execution, and other issues related to order execution. Trade throughs are therefore more common in European markets. Another important difference is that cross-market trading in European markets is more costly for investors due to larger clearing and settlement costs.

After the implementation of MiFID, most large OSE-listed stocks started trading on other stock exchanges.¹¹ As a result of market fragmentation, the OSE market share of trading in its most liquid stocks has declined from 100% in 2007 to around 40% in 2016 (Meling and Ødegaard, 2017).

2.2 Data Sources

Data for trading of US equities is collected from the Center for Securities Prices (CRSP) through the Wharton Research Data Services. We only consider ordinary shares,¹² and limit the analysis to stocks with their main listing at the NYSE.¹³ For the NYSE shares, CRSP only contains closing spreads starting at the end of 1992. Our US sample therefore starts in 1993.

The Norwegian data is from two sources. We utilize a dataset of complete trades and quotes at the OSE, provided by the market surveillance department at the OSE. The data set comprises trades for all stocks with a main listing at the OSE over the period 1999-2017. In addition the data contains all orders, and five levels of the order book at each change in the order book.

We complement the detailed trade data with market data from the Oslo Stock Exchange Data Services (OBI), which provides daily observations of stock market data, containing highs and lows, last trade price, and the closing ask and bid prices. This data source also provides corporate action data etc. The data from OBI starts in 1980 and ends in 2017.

Table 1 gives some descriptive statistics for the data sample. All measures are calculated at a monthly frequency, i.e. daily estimates are averaged over a month. Further, the time span for estimation differs across the measures. The high-low measures and the closing spread are estimated for the period 1993–2016 for the US. The corresponding Norwegian averages cover the period 1993–2017. The average of the effective spread for Norway is restricted to the period 1999–2017 where we have trade-by-trade data.

For the US, we calculate averages (and medians) for the whole exchange. We also look at averages grouped into four size groups, and three subperiods: 1993-1999, 2000-2009 and 2010–2016. We see that all measures of liquidity are improving in size, the larger stocks are more liquid. Looking at the subperiods, we see a marked fall in quoted spreads (from 2.4% to 0.18%), which is not mirrored in the

¹¹The exchanges include BATS over-the-counter, BATS CXE (former Chi-X), Turquoise, and BATS BXE (former BATS Europe)

¹²Share codes 10-12.

¹³CRSP Exchange code 1.

other liquidity measures.

For Norway, we show the same descriptives, with some additional detail. We show primarily data for stocks in the main index at the OSE, the OBX. This index contains the 25 largest and most liquid stocks at the OSE.¹⁴ The stocks in the OBX are the ones most comparable to those listed on the NYSE.¹⁵ The non-OBX firms are very small compared to the typical NYSE firm. The measures of liquidity for OSE firms show similar patterns as the NYSE, liquidity is improving with firm size, spreads fall markedly over time, a fall which is not so marked for the other liquidity measures.

Table 1: Descriptive Statistics

The tables provide summary statistics for the data. Panel A: Data for NYSE firms, 1993–2016. Panel B: Data for OSE firms, 1993–2017. For both exchanges we describe the relative (closing) spread, the two high low measures, the LOT measure, the Amihud measure, and the market capitalization. For the OSE we in addition describe the effective spread, and a measure of the typical order size. The market cap figures are in the respective currencies. The exchange rate NOK/USD at the end of 2016 was: 1 USD corresponds to 8.62 NOK.

Panel A: US (NYSE)

	NYSE firms		Size Quartiles (means)				Subperiods		
	mean	med	1 (small)	2	3	4 (large)	1993–99	2000–10	2010–16
Relative Spread (%)	1.13	0.42	1.98	1.18	0.84	0.55	2.44	0.87	0.18
High Low (CS) (%)	0.72	0.56	0.86	0.73	0.69	0.62	0.74	0.76	0.65
High Low (AR) (%)	0.85	0.67	1.01	0.87	0.81	0.74	0.89	0.91	0.74
LOT (%)	1.55	1.06	2.32	1.64	1.29	0.93	2.68	1.22	0.84
Market Cap (mill USD)	4764	752	136	511	1647	16937	2685	4889	6729
Amihud ILLIQ	3.5e-6	2.5e-9	9.8e-6	4e-7	1.3e-7	9.3e-8	7.3e-6	3.1e-6	1.7e-7

Panel B: Norway (OSE)

	OBX firms		Non-OBX firms		Size Quartiles (means)				OBX firms – Subperiods		
	mean	med	mean	med	1 (small)	2	3	4 (large)	1993–99	2000–10	2010–16
Relative Spread (%)	0.53	0.38	3.01	2.04	5.19	3.33	2.29	1.18	0.91	0.54	0.18
High Low (CS) (%)	0.97	0.82	1.98	1.41	3.23	2.17	1.62	1.08	0.80	1.14	0.88
High Low (AR) (%)	0.75	0.64	1.08	0.79	1.65	1.17	0.94	0.72	0.51	0.89	0.77
Effective Spread (%)	0.20	0.15	1.05	0.85	1.63	1.20	0.85	0.44	0.35	0.23	0.11
LOT (%)	1.77	1.58	4.04	3.12	6.15	4.29	3.40	2.32	2.31	2.03	0.93
Order Size (thousands)	28.2	8.7	21.2	0.8	4.0	3.9	21.3	44.4	32.0	29.2	23.2
Market Cap (mill NOK)	32522	10855	2100	730	184	512	1520	16627	8745	33975	51644
Amihud ILLIQ	0.0100	0.0004	1.3008	0.0851	3.8399	1.2691	0.4363	0.1022	0.0323	0.0024	0.0004

2.3 Comparing spread estimators, whole market

We start the analysis by showing the evolution of trading costs. Panel A of Figure 2 illustrates the case of NYSE by showing the cross-sectional averages of transaction cost measures, and Panel B shows corresponding averages for the largest companies at the OSE. The measures we compare are the Lesmond et al. (1999) estimator, the relative spread, the Corwin and Schultz (2012) measure and the

¹⁴The OBX is revised bi-annually, see Meling (2018) for more details about the OBX.

¹⁵The market cap of the average OBX firm is similar to the average market cap across all NYSE firms.

Amihud measure.¹⁶ The Norwegian plot also includes estimates of the effective spread for part of the period. Both market places show striking differences between liquidity measures. For the NYSE, the relative spread essentially “drops to the floor” in the 2002/2003 time frame, and stays there. This spread effect is also a result of the decimalization of the US stock markets in 2001. The Amihud measure also drops significantly around 2002 for the NYSE that could be related to inflated trading volume following market fragmentation in the US. In the case of Norway we see a more gradual fall in measured spreads and Amihud measure, but still a marked fall relative to the alternative measures.

An alternative way of illustrating properties of the transaction cost measures is to show how they vary with size. Figure 3 shows time series of liquidity measures averaged across four size-sorted portfolios at the OSE. The figure shows that the decrease in conventional spreads is most pronounced for large stocks.¹⁷ The high-low measures, however, have not changed much for any stock size. The Amihud measure has increased significantly for the smaller stocks, while it has stayed at the same level for the larger stocks.

Let us next look at the evolution of *trade size*, which we will be studying more closely. Figure 4 shows the evolution of trade size for size sorted portfolios at the OSE. We consider two estimates of trade size. One is based on daily data, where we divide the number of stocks traded during a day by the number of distinct trades during the day. This data is available for 1989–2017. The other is estimated using microstructure data, where we illustrate median transaction size during each day. This data is available for 1999–2016. Figure 4 illustrates the time series evolution of these data, broken down by firm size portfolios. This figure shows a clear tendency towards smaller transaction sizes at the OSE.


¹⁶In an internet appendix we show similar pictures for two alternative measures, the Abdi and Ranaldo (2017) estimator and the Roll (1984) estimator. These measures show the same pattern, in particular, they indicate much higher transaction costs than spreads in recent years.

¹⁷In an internet appendix we show similar pictures for size sorted portfolios for the NYSE. Here too, for the smaller stocks the recent average relative spread is larger, but still low compared to the historical estimates.

Figure 2: Time Series Spread Measures, US and Norway

The figure shows cross-sectional averages of quarterly estimates of transaction costs for NYSE companies (Panel A) for the period 1993 to 2016, and for OSE companies (Panel B) for the period 1993–2017. The OSE numbers are calculated using the 25 largest companies at the OSE (the constituents of the OBX index). Transaction costs are estimated using four methods: *LOT* is the Lesmond et al. (1999) spread estimate, *CS* is the Corwin and Schultz (2012) spread estimate, *ILLIQ* is the Amihud measure of price impact, *Spread* is the closing spread, measured as the difference between closing bid and ask divided by the mid-price. Norway also includes an estimate of the *Effective Spread*, which is the relative difference between midpoint and execution price, averaged over the day. The measures are calculated on a daily basis and averaged across days in a quarter. The averages are trimmed.

Panel A: US (NYSE) 1993–2016



Panel B: Norway (OBX firms) 1993–2017






Figure 3: Time-series evolution of spread measures: OSE – size sorted portfolios

The figures illustrates the time series development of the quarterly average of closing spread (Panel A), effective spread (Panel B), high-low spreads (Panel C and D), and Amihud measure (Panel E). The estimates are calculated separately for four size-sorted portfolios.


(a) Closing Spread




(b) Effective Spread



(c) High-Low measure (CS)



(d) High-Low measure (AR)



(e) Amihud measure





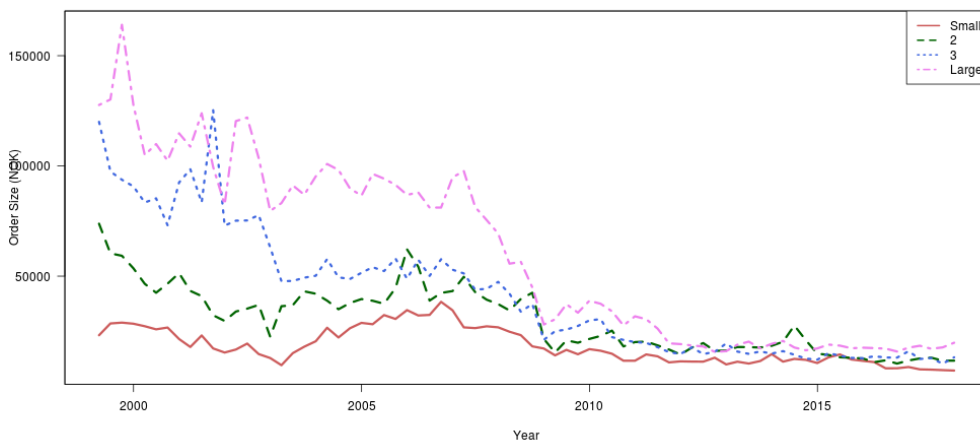
Figure 4: Time-series evolution of the average trade size

The figure illustrates time series development of the quarterly average of trade size in NOK, grouped by firm size quartile. Average across stocks in each quartile. Panel A: Trade size (in shares) is estimated as the number of shares traded in a day divided by the number of distinct trades during the day. The trade size in NOK is calculated by multiplying the number of shares with either the vwap or the close price. Panel B: Daily median trade size in NOK is calculated as the median across transactions during the day. Trimmed averages.

Panel A: Daily trade size (Volume/No Trades)



Panel B: Median trade size (calculated from transactions)



3 Does trade size matter?

What we want to investigate is to what degree transaction size affects measures of stock liquidity. To that end we employ a number of regression specifications. In addition to the variable of interest, a proxy for transaction size, we control for three additional stock characteristics: (inverse of) stock price, stock return volatility, and firm size. These characteristics have been shown to be significant determinants of transaction cost estimates in the microstructure literature.¹⁸ It is common to use the inverse of a price when explaining transaction costs since it can be regarded as a proxy for a relative tick size, which defines a lower bound for a bid-ask spread. The logarithm of market capitalization is a proxy for firm size. We expect transaction costs to be declining in market cap. More volatile stocks can be expected to be less liquid, and have higher transaction costs. As a proxy for transaction size we calculate the daily average of trade size measured in NOK.¹⁹

We estimate the effect of trade size on the transaction cost measures using five different regression specifications with the closing spread, the effective spread, the high-low spread estimators, and the Amihud measure as dependent variables. The regressions are formulated as fixed-effects panel regressions. This approach is used to control for potentially unobserved fixed effects. Specifically, we include stock-fixed effects to account for unobserved factors that vary across stocks but are invariant over time.²⁰ The analysis is restricted to shares in the OBX, the main market.²¹ The estimation results are given in Table 2.

The coefficients on control variables are consistent across regression specifications and conform to our expectations, with the inverse of a stock price and stock volatility positively related to spreads, and market cap negatively related.

For the variable of interest, transaction size, the results indicate that it has a positive and significant effect on the closing and effective spreads, which is consistent with the “measurement problem” we have discussed; that the decline in spreads is at least partly driven by the fall in transaction size. Interestingly, the high-low measures exhibit a significantly negative relation with trade size, i.e. if the high-low measures really are better estimates of transaction costs, this argues that transaction costs are increasing as trade sizes fall. If this is correct, it means trade sizes have fallen “too much”. The Amihud measure is found to be unaffected by trade size, conforming to our expectations. Therefore, the Amihud measure seems to be a robust measure of price impact costs as it does not depend on trade size.

We additionally perform regression analysis for the subperiod 1999–2007 for OBX stocks, where we can test the effect of average transaction size on liquidity before market fragmentation. The regression results for this subperiod (see panel B of Table 2) provide strong support for the order size hypothesis. The closing and effective spreads are shown to be positively affected by the average trade size, while the high-low estimators are unaffected. The Amihud measure is found to be significantly

¹⁸See for example Keim and Madhavan (1997).

¹⁹We use the OBI trade size in NOK calculated from daily data. We do not use volume in number of shares, as it depends on the stock price.

²⁰Including time-fixed effects does not change the results.

²¹In an internet appendix we also provide the regression results for the whole market. The results correspond to the results for the OBX stocks.

Table 2: Trade size effect on bid-ask spread measures: OBX stocks

The tables show the one-way fixed-effects regression results with the closing spread, the effective spread, the Amihud measure, and the high-low estimators of Corwin and Schultz (2012) (High Low (CS)) and Abdi and Rinaldo (2017) (High Low (AR)) as the dependent variables. The independent variables are the inverse of a stock price ($1/\text{Stock Price}$), stock volatility (Volatility), the logarithm of firm size ($\ln(\text{Market Cap})$), and the logarithm of trade size in NOK (Trade Size). The regressions in panel A are estimated for the OBX stocks over the period 1999–2016 at a daily frequency. The regressions in panel B are estimated for the OBX stocks over the period 1999–2007 at a daily frequency. For each model, the estimated coefficients with the Driscoll and Kraay (1998) robust standard errors (adjusted for heteroscedasticity, cross-sectional and temporal dependence) are reported, along with the regression statistics – the number of observations and Adjusted R-squared.

Panel A: 1999–2016

	<i>Dependent variable:</i>				
	Closing Spread	Effective Spread	High Low (CS)	High Low (AR)	Amihud
$1/\text{Stock Price}$	0.0060*** (0.0006)	0.0025*** (0.0002)	0.0022*** (0.0007)	0.0031*** (0.0011)	0.0739*** (0.0081)
Volatility	0.0261*** (0.0019)	0.0150*** (0.0008)	0.0498*** (0.0035)	0.0811*** (0.0053)	0.0700*** (0.0110)
$\ln(\text{Market Cap})$	-0.0055*** (0.0001)	-0.0022*** (0.00003)	-0.0008*** (0.0001)	-0.0024*** (0.0001)	-0.0311*** (0.0007)
Trade Size	0.0008*** (0.00004)	0.0003*** (0.00001)	-0.0006*** (0.00004)	-0.0003*** (0.0001)	-0.0005 (0.0004)
Observations	208,934	186,134	207,374	207,374	209,250
\bar{R}^2	0.12	0.25	0.02	0.02	0.12

Panel B: 1999–2007

	<i>Dependent variable:</i>				
	Closing Spread	Effective Spread	High Low (CS)	High Low (AR)	Amihud
$1/\text{Stock Price}$	0.0056*** (0.0013)	0.0025*** (0.0004)	0.0033*** (0.0010)	0.0043** (0.0018)	0.1327*** (0.0168)
Volatility	0.0249*** (0.0024)	0.0134*** (0.0008)	0.0279*** (0.0031)	0.0619*** (0.0059)	0.0768*** (0.0129)
$\ln(\text{Market Cap})$	-0.0055*** (0.0001)	-0.0022*** (0.00003)	-0.0010*** (0.0001)	-0.0024*** (0.0002)	-0.0370*** (0.0007)
Trade Size	0.0006*** (0.0001)	0.0002*** (0.00002)	-0.0001 (0.0001)	-0.0001 (0.0001)	0.0022*** (0.0005)
Observations	103,992	104,079	103,284	103,284	104,278
\bar{R}^2	0.11	0.23	0.01	0.01	0.15

Note:

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

positively related to the average trade size in this subsample, indicating that price impact costs were strongly associated with trade size before market fragmentation.

To mitigate an omitted variable problem that is likely to arise in modeling spreads, we suggest to look at how the average trade size affects differences between spread measures. Taking on this perspective would directly answer the question whether the changes in trade size drive the deviations between the conventional spreads and high-low measures. As shown in Table 3, the average trade size coefficient is negative and highly significant for all differences in spreads. This implies that a smaller average trade size is associated with a larger gap between the spreads and high-low measures, and vice versa.

To summarize, as also illustrated in Figure 2, the difference in levels between the conventional spreads and high-low spreads has increased significantly after 2008. During the same period, as shown in Figure 4, the average trade size dropped substantially in response to market fragmentation and algorithmic trading. This time-series observation indicates that a larger gap between spread measures could be associated with a smaller average trade size. In other words, our finding of a negative relationship between the average trade size and differences between spreads is in agreement with this time-series observation.

Table 3: Trade size effect on differences between spreads

The table shows the one-way fixed-effects regression results with the absolute differences between the closing spread and the High Low (CS) measure, the effective spread and the High Low (CS) measure, the closing spread and the High Low (AR) measure, and the effective spread and the High Low (AR) measure as the dependent variables. The independent variable is the logarithm of trade size in NOK (Trade Size). The regressions are estimated for the OBX stocks over the period 1999–2016 at a daily frequency. For each model, the estimated coefficients with the Driscoll and Kraay (1998) robust standard errors (adjusted for heteroscedasticity, cross-sectional and temporal dependence) are reported, along with the regression statistics – the number of observations and Adjusted R-squared.

	<i>Dependent variable:</i>			
	Closing Spread – High Low (CS)	Effective Spread – High Low (CS)	Closing Spread – High Low (AR)	Effective Spread – High Low (AR)
Trade Size	–0.0013*** (0.00003)	–0.0009*** (0.00002)	–0.0015*** (0.00003)	–0.0013*** (0.00003)
Observations	187,500	187,500	187,500	187,500
\bar{R}^2	0.03	0.02	0.02	0.02
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01			

4 Exogenous variation in order size

We complement the regression analysis with analysis of two cases of exogenous variation affecting transaction sizes. The regression analysis without exogenous variation in transaction size could provide biased inference due to endogeneity issues, such as omitted variable bias or simultaneity. To overcome these issues and argue more convincingly for causality it is necessary to find exogenous variation in transaction size, i.e. independent of the error term in the model. Using two quasi-experiments that provide exogenous change in transaction size for a group of stocks at the OSE, we

can estimate the causal effect of transaction size on trading costs, that is, the effect of transaction size holding all other things equal (*ceteris paribus*). We apply a difference-in-differences approach to ensure a *ceteris paribus* comparison. The quasi-experiments we use are stock splits and lot size changes, which allow us to test the effect of exogenous changes in the average transaction size and minimum order size on trading costs respectively.

4.1 Stock splits

As a first investigation of the various liquidity measures, we use stock splits as a laboratory. Stock splits offer a natural experiment to test our order size hypothesis. After a split, there are two potential effects. First, stock price decreases and consequently the relative tick size increases. Second, the average trade size decreases as a result of achieving an optimal (lower) price range and dispersed ownership represented by small investors who trade in smaller quantities. Thus, stock split events could be used for testing the effect of changes in average trade size on liquidity.

It might be useful to discuss stock splits and their effects in general first before conducting a liquidity analysis around a split. A stock split, also known as a forward split, is a corporate action of dividing existing shares into multiple shares, which directly results in a decrease of the share price and an increase of the number of shares outstanding. Stock splits do not add any market value, as the share price is adjusted accordingly (scaled by the adjustment factor), i.e. an increase in the number of shares is offset by a decrease in the share price. It is best illustrated by an example. In a 2-for-1 split, an additional share is given for each share held by an investor, but the share value is reduced by half (the adjustment factor is 0.50). Hence, the market value remains the same after the split.

Empirical research has shown that stock splits are followed by an increase in volatility and proportional spreads. It remains a puzzle why firms split their stocks. Several explanations for stock splits have been suggested in the literature. A classical argument is the trading range hypothesis of Copeland (1979). In this view, firms may want to keep their stock price within a lower price range in order to attract new clientele such as uninformed or small investors, or achieve greater dispersion of ownership. A smaller average trade size and an enlarged ownership base typically reported after the split provide support for this hypothesis. However, liquidity is shown to decrease after the split, opposed to what the trading range hypothesis predicts. Another hypothesis is related to a reduction in information asymmetry. Stock splits are commonly regarded as a costly signal of “good information.” Stock splits also attract attention and the split factor is shown to be positively related to analyst following. Desai, Nimalendran, and Venkataraman (1998), however, find that adverse selection increases after the split, which contradicts the prediction regarding lower information asymmetry. An increase in spreads following the split is still consistent with the signaling story, though. Finally, an optimal tick size hypothesis, proposed by Harris (1996) and Angel (1997), suggests that firms split their stocks to increase a relative tick size and spreads as a result. According to this hypothesis, higher spreads induce more liquidity provision. For instance, uninformed investors would shift from market orders to limit orders as the latter become more profitable. An increase in spreads after the split is consistent with this hypothesis, however, it remains unclear whether overall liquidity improves due to enhanced liquidity provision.

The effect of stock splits on liquidity is a controversial question due to inconsistent evidence provided in the empirical studies. The literature has reported both increases and decreases in the bid-ask spreads following stock splits. There are also conflicting predictions from the trading range and signaling hypotheses. A trading range hypothesis predicts a positive effect of stock splits on liquidity as the lower share price should attract more investors. A signaling hypothesis, on the contrary, predicts a negative effect of stock splits on liquidity as they are regarded as a costly signal of positive information. The latter hypothesis seems to have gained more empirical support in the literature.

In this paper, however, we hypothesize that stock splits do not convey any information about a firm, and have purely mechanical effects, such as effects on the share price and the number of shares outstanding. Consequently, the changes in transaction costs after the split are directly related to the changes in the share price and the average trade size as the number of shares in an order (trade) changes. Thus, the stock split event gives us an opportunity to study the effect of the exogenous shocks to the stock price (and a relative tick size, which is an inverse of the stock price) and the average trade size on transaction costs. A negative relationship between the stock price and proportional spreads has been documented in many studies. There is also evidence of a negative effect of the stock price on proportional spreads in the stock split event study of Conroy, Harris, and Benet (1990). Hence, a decrease in the stock price after a split should lead to higher transaction costs. Most studies also report a decrease in the average trade size after stock splits (Ferris, Hwang, and Sarin, 1995; Guo, Zhou, and Cai, 2008), which implies that shares are distributed to many small investors who trade in small quantities. Further, Desai et al. (1998) shows that adverse selection costs become higher after the split as the proportion of informed trading increases. The reason for this could be greater analyst coverage accompanying stock splits. Informed traders trade in small quantities to hide their trades among small trades by uninformed investors. Though, as pointed out by Easley, O'Hara, and Saar (2001), the proportion of uninformed trading also increases, which should lead to lower adverse selection costs. Therefore, the effect on spreads is unclear as the net adverse selection costs could either increase or decrease. As shown in Easley et al. (2001), increased trading by both uninformed and informed investors lead to a larger number of trades, and higher volatility as a result. Hence, an increase in spreads is also attributed to an increase in volatility after the split. This evidence suggests that lower stock price and smaller average trade size should be associated with higher spreads after the split.

4.1.1 OSE sample of stock splits


The frequency distributions of stock splits and split types in our sample are presented in Figure 5. The largest number of splits was reported in the year 2000. Since then the number of stock splits has been decreasing. The stock splits with split factors 2-for-1, 4-for-1, 5-for-1, and 10-for-1 are the most frequent in our sample. From Table 4 we can see that the median trade size decreases substantially after the split along with the share price, while return volatility increases for splitting stocks. The observed changes in stock characteristics are consistent with our expectations and previous literature. Non-splitting stocks, however, have similar characteristics before and after the split. This implies that

the changes in stock characteristics for splitting stocks were driven by the stock split alone. However, there seem to be inherent differences between splitting and non-splitting stocks. By design, splitting stocks are high-priced stocks (232 NOK), while non-splitting stocks are stocks with a fairly average price (55 NOK). This difference in prices, however, is mostly eliminated after the split. Hence, the results from comparison of splitting stocks and non-splitting stocks before and after the split might be affected by the potential differences between these two groups of stocks.


Figure 5: Frequency distribution of stock splits

The figure illustrates the frequency distributions of stock splits by year (Panel A) and by split type (Panel B) for the period 1999–2015. The sample consists of 94 stock splits in the OSE stocks.

Panel A: Frequency of splits by year



Panel B: Frequency of splits by type



4.1.2 Analysis of liquidity before and after a split

We use split dates²² from 94 stock splits at the OSE during the period 1999–2014 and compare the crosssectional averages of spreads before and after the split. To isolate the effect of the split, a short event window of 60 trading days is used (30 days before the split and 30 days after the split respectively). Furthermore, we conduct the same analysis for a random control sample of non-splitting stocks, in order to ensure that liquidity effects observed for splitting stocks around the split dates are not caused by the changes in macroeconomic activity.

Table 4: Stock characteristics before and after a split

The table shows the crosssectional median of stock characteristics over the 5-day period before and the 5-day period after the split date for both splitting and non-splitting stocks. Stock price is given in NOK.

Variable	Before a split	After a split
Panel A: Splitting stocks		
Stock price	231.8	84.9
Stock volatility	0.005	0.007
Trade size (NOK)	52617.1	7163.7
Trade size (shares)	229.3	84.7
Panel B: Nonsplitting stocks		
Stock price	54.9	55.1
Stock volatility	0.006	0.006
Trade size (NOK)	3233.4	2653.4
Trade size (shares)	56.7	51.0

In our analysis, we not only look at the benchmark spreads, but also at the high-low measures and the Amihud measure around the split dates. Based on the literature, we expect spreads to increase significantly after the split. However, we posit that the high-low measures and the Amihud measure will be less affected, as they should be robust to changes in the average trade size.

Figure 6 shows the time-series development of spread measures, such as the closing spread, the effective spread, and the high-low measures, and the Amihud measure within a 60-day event window. As shown in Figure 6, all measures increase significantly after the split. This behavior could be attributed to the joint effect of the stock price and the average trade size.


Figure 7 additionally shows liquidity reactions to stock splits for randomly selected non-splitting stocks. As revealed by the plots, there are no significant changes in spread estimates for non-splitting stocks around the split date. However, the closing spread, the Amihud measure, and high-low spread estimates show an increase after the split, which indicates that the increase in transaction costs for splitting stocks could be driven by general macroeconomic changes. The effective spread estimates for non-splitting stocks remain unchanged after the split. This implies that the increase in effective spreads for splitting stocks could be a direct consequence of the split.

²²The analysis is centered around the ex-date. The announcement date is typically very close to the ex-date for stocks in our sample and therefore not considered in the analysis. Furthermore, we assume that stock splits have no information content.


Figure 6: Analysis of liquidity before and after a split: splitting stocks

The figure illustrates time-series development of spread measures for splitting stocks around the stock split date. The event window of 60 trading days is used. The average spread before the split is compared to the average spread after the split using the paired two sample t-test for differences in means. The mean values of spreads in two periods (μ_1 and μ_2) and t-statistics are reported under the time-series plots.

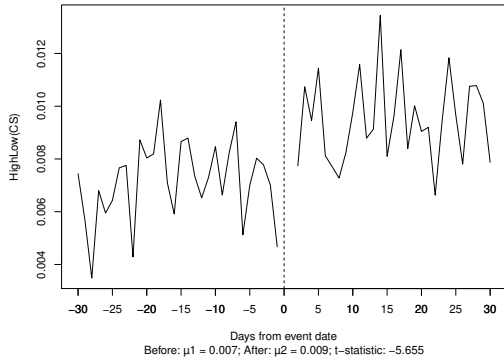
(a) Closing spread




(b) Effective spread



(c) High-Low measure (CS)



(d) High-Low measure (AR)



(e) Amihud measure






Figure 7: Analysis of liquidity before and after a split: non-splitting stocks

The figure illustrates time-series development of spread measures for non-splitting stocks around the stock split date. The event window of 60 trading days is used. The average spread before the split is compared to the average spread after the split using the two sample t-test for differences in means. The mean values of spreads in two periods (μ_1 and μ_2) and t-statistics are reported under the time-series plots.


(a) Closing spread




(b) Effective spread




(c) High-Low measure (CS)



(d) High-Low measure (AR)



(e) Amihud measure



4.1.3 Difference-in-differences analysis

The liquidity comparison before and after a split is not sufficient for claiming causality due to potential inherent differences between splitting and non-splitting stocks and changes in general market conditions that affect both groups. One way to control for these differences is to perform a difference-in-differences analysis. We estimate the following regression to infer the causal effect²³ of stock splits on the spread measures.

$$y_{it} = \beta TREAT_i + \gamma POST_t + \delta(TREAT_i * POST_t) + \varepsilon_{it}, \quad (5)$$

where $TREAT_i$ is a dummy for the treatment stock, $POST_t$ is a dummy for the post-treatment period, and $TREAT_i * POST_t$ is an interaction term. The treatment dummy controls for fixed differences between splitting and non-splitting stocks and indicates stocks in the treatment group, i.e. splitting stocks. The post-treatment period dummy controls for the change in general conditions over time and indicates the period from the split date onward. Lastly, the interaction term between these two dummies indicates the splitting stocks after the split.

We perform the two-way fixed-effects difference-in-differences estimation, which leads to elimination of the treatment dummy, however, the post-treatment period dummy and the interaction term remain. The coefficient of interest here is the coefficient of the interaction term (δ), which is the causal effect of stock split on spreads. Table 5 presents the regression results from the difference-in-differences estimation. Cluster-robust standard errors are used to adjust for potential heteroscedasticity and serial correlation. The regression results indicate that there is a positive and significant causal effect of stock split on the effective spread. This finding contradicts the order size hypothesis as it implies a negative relationship between the effective spread and average trade size, i.e. the average trade size decreases after the split while the effective spread increases. Though, the stock split is also associated with an increase in a relative tick size and stock volatility. This could explain an increase in the effective spread estimates following the split. Hence, it does not necessarily reflect the effect of the average trade size. Consistent with the order size hypothesis, we find no significant effect on the closing spread, the Amihud measure, and the high-low estimators. This implies that these measures are unaffected by the changes in the average trade size and relative tick size following the split.

4.2 Lot size revisions

To disentangle the effect of order size on liquidity from the joint effect of stock price and order size, we conduct a liquidity analysis around lot size revisions at the OSE. Lot size is the number of shares in a transaction, e.g. lot size 1, which is a standard for all equities at the OSE since migration to TradElect trading system in April 2010, means that the minimum order size is one share. Prior to standardization in 2010, lot sizes for a subset of OSE stocks were revised by the stock exchange on a quarterly basis. The point of departure for lot size assignment was that the value of the position should be around 10000 NOK for stocks listed on the Main List, and around 5000 NOK for stocks listed on the SMB List (small- and medium-sized companies). The lot sizes for different stocks

²³Given that the parallel trends assumption is satisfied.

Table 5: Difference-in-differences results: stock splits

The table shows the estimates of treatment effect coefficients from the two-way fixed-effects difference-in-differences estimation with the closing spread, the effective spread, the Amihud measure, and the high-low estimators of Corwin and Schultz (2012) and Abdi and Rinaldo (2017) as the dependent variables. A 60-day event window is used in the estimation. For each model, the estimated coefficients with the clustered (at the stock level) robust standard errors are reported, along with the regression statistics – the number of observations and Adjusted R-squared.

	<i>Dependent variable:</i>				
	Closing Spread	Effective Spread	High Low (CS)	High Low (AR)	Amihud
POST	0.0003 (0.0021)	0.0004 (0.0003)	-0.0009 (0.0008)	-0.0009 (0.0015)	0.0881** (0.0433)
TREAT*POST	0.0003 (0.0018)	0.0009*** (0.0003)	0.0014 (0.0010)	0.0014 (0.0014)	-0.0457 (0.0553)
Observations	7,232	5,002	5,319	5,319	5,258
\bar{R}^2	-0.39	-0.53	-0.53	-0.53	-0.36

Note:

*p<0.1; **p<0.05; ***p<0.01

were revised accordingly to account for their respective price changes, i.e. a lot size was adjusted downward when a stock's price was rising and upward when a stock's price was falling. Figure 8 shows a frequency distribution of lot sizes for stocks subject to quarterly lot size revisions. The most typical lot sizes are 500, 1000, 200, 2000, and 100 shares. This figure also indicates that the majority of lot size revisions, both decreases and increases in lot sizes, were carried out in non-OBX stocks.

We use this natural experiment to test the effect of minimum order size changes on various liquidity measures.²⁴ We employ a difference-in-differences approach to estimate the causal effect of minimum order size changes on liquidity. We expect the changes in minimum order size to have similar effects as the changes in average order size, i.e. the closing and effective spreads should be significantly positively affected, whereas the high-low estimators and the Amihud measure should not be affected by these changes.

Our data set comprises quarterly data on lot size revisions for a number of OSE stocks during 2002-2010²⁵. We filter the data using several criteria. First, we separate lot size increases and lot size decreases in order to test for possible asymmetric effects of order size changes on liquidity. Second, we restrict both samples to sufficiently large changes in lot size, where a lower bound is set arbitrarily to 3000 shares. Smaller lot size changes are less likely to cause price changes and therefore would not have any significant effect on any of the spread measures. Finally, we remove lot size revisions that took place in June-October 2009 as those revisions overlap with minimum tick size changes for OBX stocks at the OSE and could potentially bias the results. We also ensure that the estimation windows for lot size increases and decreases do not overlap for the same stocks. For that reason we use a relatively short estimation window of 60 trading days (30 days before and after the revision). The final samples consist of 66 lot size decreases and 98 lot size increases. We additionally collect


²⁴In an internet appendix we provide event study plots that show liquidity changes around lot size decreases and increases for both revised and unrevised stocks.

²⁵This data set was collected by Tom Meling from the lot size announcements on NewsWeb. We are grateful to Tom for providing the data.


Figure 8: Frequency distribution of lot sizes and stock types

The figure illustrates the frequency distribution of lot sizes (Panel A) and proportion of OBX vs. non-OBX stocks that were subject to lot size decreases and increases (Panel B) during the period 2002–2009.

Panel A: Frequency distribution of lot sizes



Panel B: Frequency distribution of OBX/non OBX stocks by revision type




matched control samples of stocks not seeing lot size revisions during the sample period, but still being listed during the corresponding revision periods.

The frequency distributions of lot size revisions by year and type are given in Figure 9. As illustrated in this figure, lot size decreases have been more frequent than lot size increases during 2003–2007, however, lot size increases outnumbered lot size decreases significantly in 2002, and 2008–2009. The largest number of revisions, which were mainly lot size increases, followed the financial crisis in 2008. The frequency distribution of lot size revisions by type indicates that the most frequent absolute changes in lot sizes have been 3000, 5000, and 8000 shares.


Figure 9: Frequency distribution of lot size revisions

The figure illustrates the frequency distributions of lot size revisions by year (Panel A) and by the absolute change in lot size (Panel B) for the period 2002–2009. Two samples of lot size decreases and increases consist of 66 and 98 lot size revisions respectively.

Panel A: Frequency of lot size revisions by year



Panel B: Frequency of lot size revisions by type



We implement the fixed-effects difference-in-differences estimation with both stock- and date-fixed effects to account for unobserved effects invariant across time and stocks respectively. Including stock-fixed effects leads to elimination of the TREAT dummy, which signifies stocks exposed to treat-

ment, i.e. lot size revisions. The POST dummy, which stands for the post-treatment period, however, is not omitted when we add date-fixed effects as we have variation in treatment timing. The coefficient on TREAT*POST is an estimated causal effect of lot size changes on liquidity measures. The treatment effect coefficient estimates for lot size decreases and increases are reported in panels A and B in Table 6.

The results indicate that lot size decreases have a significant negative effect on the effective spread, which means that effective spreads fall in response to the fall in minimum order size. The decrease in minimum order size, however, has no significant effect on the closing spread. The high-low estimators are shown to be unaffected by the lot size decrease, which implies that the high-low estimates should be less downward-biased in the case of split orders.

The lot size increase is found to have similar effects on the liquidity measures. Specifically, it has a significant positive effect on the effective spreads and no significant effect on the high-low estimators and the Amihud measure. The closing spread, however, is shown to be significantly positively affected by the increase in minimum order size. This result is in line with our expectations, as the minimum order size sets a lower bound on the order size and therefore its increase will result in higher quoted (closing) spreads (due to less favorable quotes for larger quantities).

The results for lot size revisions are therefore consistent with the order size hypothesis. Furthermore, they suggest that the high-low estimators as well as the Amihud measure are more robust to both order size decreases and increases than the conventional spreads.

Table 6: Difference-in-differences results

The table shows the estimates of treatment effect coefficients from the two-way fixed-effects difference-in-differences estimation with the closing spread, the effective spread, the Amihud measure, and the high-low estimators of Corwin and Schultz (2012) and Abdi and Ranaldo (2017) as the dependent variables. A 60-day event window is used in the estimation. For each model, the estimated coefficients with the clustered (at the stock level) robust standard errors are reported, along with the regression statistics: the number of observations and Adjusted R-squared.

Panel A: Lot size decreases

	<i>Dependent variable:</i>				
	Closing Spread	Effective Spread	High Low (CS)	High Low (AR)	Amihud
POST	-0.0090 (0.0083)	0.0005 (0.0006)	-0.0145** (0.0074)	-0.0181 (0.0123)	0.1123*** (0.0396)
TREAT*POST	-0.0032 (0.0034)	-0.0011** (0.0005)	-0.0040 (0.0031)	-0.0055 (0.0034)	-0.0483 (0.0323)
Observations	6,033	4,185	4,405	4,405	5,258
\bar{R}^2	-0.29	-0.46	-0.41	-0.42	-0.36

Panel B: Lot size increases

	<i>Dependent variable:</i>				
	Closing Spread	Effective Spread	High Low (CS)	High Low (AR)	Amihud
POST	-0.0054 (0.0096)	-0.0007 (0.0010)	0.0006 (0.0050)	-0.0028 (0.0061)	-0.0068 (0.1347)
TREAT*POST	0.0076** (0.0035)	0.0009* (0.0005)	0.0007 (0.0017)	0.0049 (0.0031)	-0.0538 (0.0463)
Observations	8,732	4,893	5,687	5,687	7,227
\bar{R}^2	-0.14	-0.25	-0.22	-0.22	-0.17

Note:

*p<0.1; **p<0.05; ***p<0.01

5 Conclusion

Have the institutional and technological changes to equity markets improved market quality for traders at the exchanges? To answer this question we need time series of liquidity measures invariant to changes just affecting the inner working of intermediaries. In this paper we look at one aspect of such invariance, the search for a liquidity measure that is invariant to changes in trade size.

We consider a number of liquidity measures, such as the closing and effective spreads, the Amihud measure, and the high-low estimators of Corwin and Schultz (2012) and Abdi and Rinaldo (2017). We show the evolution of these measures for two market places, NYSE and OSE. The spreads and the Amihud measure have fallen significantly on both exchanges following market fragmentation and introduction of algorithmic trading. The high-low spread estimators, however, have not fallen.

We argue that the fall in conventional spreads is at least partially driven by the fall in average order size. To the extent spread measures include price impact, a smaller average order size should be associated with lower estimated spread. The high-low estimators of Corwin and Schultz (2012) and Abdi and Rinaldo (2017) are less likely to be affected by the average trade size.

We test our hypothesis regarding the effect of order size on the various transaction cost measures using daily transaction data for the OBX stocks over the period 1999–2016. A number of fixed-effects panel regressions are estimated to determine the effect of the average trade size on the liquidity measures. The results provide strong evidence in favor of the order size hypothesis. We find that the average trade size significantly positively affects the closing and effective spreads, which confirms our notion that the fall in spreads is driven by the fall in order sizes. Interestingly, the average trade size is shown to have a significant negative effect on the high-low estimators when we use the whole period 1999–2016 in the estimation. The effect of average trade size on the high-low estimators is however insignificant for the subperiod 1999–2007, which is the period before market fragmentation at the OSE. This indicates that the high-low estimators are not tied to the typical order size, contrary to spreads. The Amihud measure is shown to be significantly positively affected by trade size for the subperiod 1999–2007, but unaffected for the whole sample period.

To overcome potential endogeneity issues in the panel regression analysis, we additionally look at liquidity changes around stock splits (corporate events that provide exogenous shocks to the average trade size and relative tick size). The difference-in-differences analysis of stock splits indicates that the high-low measures and the Amihud measure are unaffected by changes in average trade size and relative tick size after the split, while the benchmark spreads show a significant increase.

To test our hypothesis in a more direct way, we use minimum order size revisions at the OSE and perform a difference-in-differences analysis to obtain the causal effect of order size changes on liquidity measures. The results suggest that the closing and effective spreads are positively affected by the minimum order size, while the high-low estimators and the Amihud measure are robust to both minimum order size increases and decreases.

An important implication of these findings is that closing and effective spreads are prone to underestimate transaction costs of split orders, whereas the high-low estimators and the Amihud measure are not as sensitive to the fall in trade sizes. The Amihud measure, however, may still underestimate price impact. This bias can be driven by inflated trading volume as a result of increased

trading between financial intermediaries. Thus, the high-low estimators should be preferred to both conventional spread measures and volume-based measures in modern fragmented markets.

It is noteworthy, however, that despite the high-low estimators being able to remedy the problem with measurement of price impact, they could still underestimate total transaction costs of split orders as they fail to capture opportunity costs due to non-execution. The implementation shortfall would therefore be a better measure of transaction costs of split orders as it takes into account time dimension of liquidity. In practice, however, it is difficult to obtain data on split orders that allows estimation of implementation shortfall. Therefore, the high-low estimators present a good alternative since they are less downward biased than the conventional spread measures and can be easily estimated from daily high and low prices.

This result leads to our argument that equity transaction costs have not gone down in recent years, contrary to the common belief that liquidity has improved substantially in the new high frequency world.

Let us close by pointing out some regulatory implications of our results. What we show is that *order size matters*, and one should evaluate whether the very small trade sizes in today's electronic market places are too low. There are obvious ways for exchanges to encourage traders to split orders less. The lowest hanging fruit is the exchange fees. For most exchanges, fees have fixed and variable components. If one wants to encourage larger trade sizes, the fixed component can be increased and the variable component decreased. Another obvious possibility is a direct regulation establishing a minimum order size. There are other regulatory interventions that either are being implemented, or being suggested, which will also affect incentives for order splitting.

References

- Farshid Abdi and Angelo Ranaldo. A simple estimation of bid-ask spreads from daily close, high, and low prices. *The Review of Financial Studies*, 30(12):4437–4480, 2017.
- Yakov Amihud. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56, 2002.
- James J Angel. Tick size, share prices, and stock splits. *The Journal of Finance*, 52(2):655–681, 1997.
- James J Angel, Lawrence Harris, and Chester S Spatt. Equity trading in the 21st century. *Quarterly Journal of Finance*, 1(1):1–53, 2011.
- Yashar H Barardehi, Dan Bernhardt, and Ryan J Davies. Trade-time measures of liquidity. *The Review of Financial Studies*, 32(1):126–179, 03 2018.
- Michael Bleaney and Zhiyong Li. The performance of bid-ask spread estimators under less than ideal conditions. *Studies in Economics and Finance*, 32(1):98–127, 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.
- Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Recent trends in trading activity and market quality. *Journal of Financial Economics*, 101:243–263, 2011.
- Robert M Conroy, Robert S Harris, and Bruce A Benet. The effects of stock splits on bid-ask spreads. *The Journal of Finance*, 45(4):1285–1295, 1990.
- Thomas E Copeland. Liquidity changes following stock splits. *The Journal of Finance*, 34(1):115–141, 1979.
- Shane A Corwin. Dealing with negative values in the high-low spread estimator, a comment on. Working Paper, downloaded from Corwin’s homepage, 2014.
- Shane A Corwin and Paul Schultz. A simple way to estimate bid-ask spreads from daily high and low prices. *The Journal of Finance*, 67(2):719–760, 2012.
- Anand S Desai, Mahendrarajah Nimalendran, and Subu Venkataraman. Changes in trading activity following stock splits and their effect on volatility and the adverse-information component of the bid-ask spread. *Journal of Financial Research*, 21(2):159–183, 1998.
- John C Driscoll and Aart C Kraay. Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics*, 80(4):549–560, 1998.
- David Easley, Maureen O’Hara, and Gideon Saar. How stock splits affect trading: A microstructure approach. *Journal of Financial and Quantitative Analysis*, 36(1):25–51, 2001.
- David Easley, Marcos Lopez de Prado, and Maureen O’Hara. Discerning information from trade data. *Journal of Financial Economics*, 120(2):269–285, 2016.
- Gregory W Eaton, Paul J Irvine, and Tingting Liu. Measuring institutional trading costs and the implications for finance research: The case of tick size reductions. Working Paper, SSRN, November 2018.
- Stephen P Ferris, Chuan-Yang Hwang, and Atulya Sarin. A microstructure examination of trading activity following stock splits. *Review of Quantitative Finance and Accounting*, 5(1):27–41, 1995.
- Kingsley Fong, Craig W Holden, and Charles A Trzcinka. What are the best liquidity proxies for global research? *Review of Finance*, 2017.
- Thierry Foucault, Marco Pagano, and Ailsa Röell. *Market liquidity: Theory, evidence, and policy*. Oxford University Press, 2013.

- Andrea Frazzini, Ronen Israel, and Tobias Moskowitz. Trading costs. Working Paper, AQR, April 2018.
- Fang Guo, Kaiguo Zhou, and Jinghan Cai. Stock splits, liquidity, and information asymmetry – an empirical study on Tokyo stock exchange. *Journal of the Japanese and International Economies*, 22(3):417–438, 2008.
- Lawrence Harris. Statistical properties of the roll serial covariance bid/ask spread estimator. *The Journal of Finance*, 45(2):579–590, 1990.
- Lawrence Harris. Does a large minimum price variation encourage order exposure? Working Paper, Univ. of Southern California, 1996.
- Joel Hasbrouck. Liquidity in the futures pits: Inferring market dynamics from incomplete data. *Journal of Financial and Quantitative Analysis*, 39(2):305–326, 2004.
- Craig W Holden. New low-frequency spread measures. *Journal of Financial Markets*, 12(4):778–813, 2009.
- Craig W Holden and Stacey Jacobsen. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. *Journal of Finance*, LXIX(4), August 2014.
- Craig W Holden, Stacey Jacobsen, and Avanidhar Subrahmanyam. The empirical analysis of liquidity. *Foundations and Trends in Finance*, 8(4):263–365, 2013.
- Roger D Huang and Hans R Stoll. The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4):995–1034, 1997.
- Kjell Jørgensen, Johannes A Skjeltorp, and Bernt Arne Ødegaard. Throttling hyperactive robots - Order to trade ratios at the Oslo Stock Exchange. *Journal of Financial Markets*, 37(1):1–16, 2018.
- Nina Karnaukh, Angelo Ranaldo, and Paul Söderlind. Understanding fx liquidity. *Review of Financial Studies*, 28(11):3073–3108, 2015.
- Donald B Keim and Ananth Madhavan. Transactions costs and investment style: an inter-exchange analysis of institutional equity trades. *Journal of Financial Economics*, 46(3):265–292, 1997.
- Donald B Keim and Ananth Madhavan. The cost of institutional equity trades. *Financial Analysts Journal*, 54(4):50–69, 1998.
- Albert S Kyle. Continuous auctions and insider trading. *Econometrica*, pages 1315–1335, 1985.
- David A Lesmond, Joseph P Ogden, and Charles A Trzcinka. A new estimate of transaction costs. *Review of Financial Studies*, 12(5):1113–1141, 1999.
- Chien-Chih Lin. Estimation accuracy of high–low spread estimator. *Finance Research Letters*, 11(1):54–62, 2014.
- Paolo Mazza. Rethinking zero returns in the liquidity puzzle of a limit order market. *Finance*, 36:7–36, 2015.
- Tom Meling and Bernt Arne Ødegaard. Tick size wars, high frequency trading, and market quality. Working paper, University of Stavanger, 2017.
- Tom Grimstvedt Meling. Anonymous trading in equities. Working Paper, SSRN, June 2018.
- Albert J Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.
- 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 Skjeltorp, and Bernt Arne Ødegaard. Liquidity at the Oslo Stock Exchange. May 2008. Working Paper Series, Norges Bank, ANO 2008/9.
- Maureen O’Hara. High frequency market microstructure. *Journal of Financial Economics*, 116(2):257 – 270, 2015.

Richard Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39(4):1127–1139, 1984.

Raphael Schestag, Philipp Schuster, and Marliese Uhrig-Homburg. Measuring liquidity in bond markets. *Review of Financial Studies*, 29(5):1170–1219, 2016.