EFA
Eastern
Finance
Association

The Financial Review 49 (2014) 345–369

The
Financial
Review

High-Frequency Trading and the Execution Costs of Institutional Investors

Jonathan Brogaard*

Foster School of Business, University of Washington

Terrence Hendershott

Haas School of Business, University of California Berkeley

Stefan Hunt

Financial Conduct Authority

Carla Ysusi

Financial Conduct Authority

Abstract

This paper studies whether high-frequency trading (HFT) increases the execution costs of institutional investors. We use technology upgrades that lower the latency of the London Stock Exchange to obtain variation in the level of HFT over time. Following upgrades, the level of HFT increases. Around these shocks to HFT institutional traders' costs remain unchanged. We find no clear evidence that HFT impacts institutional execution costs.

**Corresponding author*: Foster School of Business, University of Washington, 434 Paccar Hall, Box 353226, Seattle, WA 98195; Phone: (206) 685-7822; Fax: (206) 513-7472; E-mail: brogaard@uw.edu.

The views in this paper are solely those of the authors and do not reflect the views of the Financial Conduct Authority (FCA) or any of the trading venues or market participants mentioned in the paper. All errors or omissions are the authors' sole responsibility. No stock exchange is party to this research or responsible for the views expressed in it. References in this report to data from the exchanges refer to information provided by the exchange to the FCA at the FCA's request. The data provided by the exchange to the FCA are anonymous. All data are aggregated and specific member firms and their activities cannot be identified. We thank the anonymous referees and the editor for helpful comments.

Keywords: high frequency trading, execution costs, institutional investors

JEL Classifications: G10, G23, O16

> *Institutional investors also have expressed serious reservations about the current equity market structure. . . . [I]nstitutional investors questioned whether our market structure meets their need to trade efficiently and fairly, in large size.*
>
> Mary Schapiro, SEC Chairman, September 7, 2010 speech.

1. Introduction

Transaction costs matter.[1] Financial markets exist so investors can efficiently transfer assets and their associated payoffs and risks; the cheaper it is to transfer an asset, the more likely the most-suited investor will end up holding the asset. In addition, with lower transaction costs, investors with private information can more readily buy and sell, aiding price discovery.

High-frequency trading (HFT) accounts for an increasingly large fraction of financial market trading, potentially affecting transaction costs. HFT is a subset of computer-based trading, defined by the use of sophisticated trading algorithms and the ability to trade rapidly to generate returns. Until recently, human intermediaries, such as NYSE Specialists and registered market makers, facilitated the smooth transfer of assets. Now, many human market makers have been substituted by computers.

While the rise of machines has raised concern, most academic evidence suggests it has improved measures of market quality such as volatility, price discovery, and liquidity. For example, Menkveld (2013) studies the entrant of a new HFT firm, and finds that after the HFT firm enters, spreads decrease by 50%. Malinova, Park and Riordan (2013) show that an increase in order submission fees results in a decrease in order submissions, and therefore the cost of trading for retail investors increases. Carrion (2013) shows that HFT firms provide more liquidity when spreads are wide. Brogaard, Hendershott and Riordan (2013) find that HFT contributes to price discovery. We contribute to the literature by examining the link between HFT activity and institutional trading costs.

Even though market-wide measures of market liquidity may improve, including the spread, this does not necessitate that institutional investors are better off.[2] The improvement in market quality may be only for a select group of market participants.

[1] We use the term "transaction costs" to mean all costs incurred in financial trading, including execution cost, commissions and rebates, information technology costs and other costs. We use the term "execution costs"—synonymously, trading costs—to mean market-adjusted execution shortfall, the volume-weighted percentage difference between the price available in the market when brokers receive institutional orders and the price at which the order is executed.

[2] Institutional investors refer to buy–side institutions such as pension plans and money managers. Our data come from Abel Noser, a well-known consulting firm that works with pension plan sponsors and money managers to monitor their equity trading costs.

Some claim that execution costs, a component of transaction costs, could be increasing because of HFT. Possible reasons include faster reaction to public information by high-frequency traders (HFTs), which could allow HFTs to pick off orders from slower market participants, or trading in front of institutional investors through the detection of autocorrelation in order flow caused by institutional investors entering large trades. Important gaps exist in the literature on the impact HFT has on the different components of the transaction costs of institutional investors including execution costs.

This paper aims to address one of these gaps. We construct measures of HFT activity and institutional investor execution costs. We show that HFT activity increases following improvements in exchange speed. From 2007 to 2011, the London Stock Exchange (LSE) implemented a variety of improvements to its technology that dramatically increased exchange speed. Using these changes in exchange speed, we study the role of HFT in institutional execution costs. We find no relationship between these shocks to the activity of HFTs and institutional execution costs.[3]

We focus on the market-adjusted execution shortfall costs as the measure of interest with respect to execution costs. The market-adjusted execution shortfall is widely used by academics as well as practitioners (Anand, Irvine, Puckett and Venkataraman, 2012, 2013). Using the price at the time the institution decides to trade as the "true" price minimizes the effect of temporary price pressures.

We study institutional execution costs in the largest 250 U.K.-listed stocks using the Abel Noser data set and show that costs have been decreasing since 2003, albeit with an interruption during the financial crisis. This time trend is consistent with work done using U.S. data (Anand, Irvine, Puckett and Venkataraman, 2013). While the data allow us to study the execution component of trading costs, we are unable to examine institutions' total trading costs. For instance, we lack data on the costs incurred by the firm in buying or developing algorithms to enter orders.

From November 2007 to August 2011, we observe HFT activity using the Financial Services Authority's (FSA) Sabre II data set.[4] The data set includes all transactions by observable HFTs in the largest 250 U.K.-listed stocks.[5] Observable HFTs in the FSA data set are those that are either directly regulated in the European Economic Area (EEA) or trade through a broker. Using a further data set from the three largest trading venues in the United Kingdom, we show that Sabre II captures

[3] As noted, this is only one component of their transaction costs. We do not study, for example, whether HFTs have increased commission costs by increasing the number of trades to fill an order.

[4] Sabre II covers all transactions of EEA regulated firms in all debt, equity and debt and equity derivative instruments listed in the United Kingdom and so it is a rich record of trading in one of the world's major financial centers. It has, however, only been used in two previous research papers (Gondat-Larralde and James, 2008; Benos and Sagade, 2012). As of April 1, 2013, the FSA no longer exists and one of its successor organizations, the Financial Conduct Authority, maintains the transaction record data.

[5] All brokers are regulated and must report the transactions of their clients. We do not observe the trades of unregulated HFT firms that are placed directly on trading venues.

70–80% of HFT activity through July 2010.[6] We therefore focus our main analysis on the period from November 2007 to July 2010.

To study the role of HFT in institutional execution costs, we regress HFT activity (at the stock-day level) on the speed of the LSE system, controlling for long-term trends in HFT activity to isolate the short-run impact of the infrastructure upgrades. For two of the four LSE system changes before August 2010, we find that HFT activity increases after the speed increase. We find no measurable change in execution costs around the technology improvements. However, one cannot draw causal inferences from this analysis. There may be a latent factor influencing both HFT activity and executions, or the direction of causality may be reversed: execution costs may influence HFT activity.

To overcome the potential endogeneity we implement a two-stage least-squares (2SLS) regression methodology. In stage one HFT activity is regressed on the instrument and control variables to find the relationship between the instrument and HFT activity. In the second stage, we use the estimated instrument coefficient from the first stage to isolate the component of HFT activity that is not driven by execution costs. The instrumental variable used to isolate this element of HFT activity is the speed change in the LSE's matching engine. The model is estimated for 20-day windows around the four TradElect upgrades separately. We fail to find an effect of HFT activity on institutional execution costs.

As one cannot prove a null hypothesis we are unable to establish that an increase HFT from current base levels does not affect institutional execution costs. However, we fail to conclude that HFTs affect institutional execution costs. There are limitations to our analysis. For instance, intraday prices are noisy, which makes execution cost measures have high variance, making small changes in execution costs difficult to detect; and the changes in the level of HFT we find are relatively small. In addition, we are examining a market that already has a high level of HFT. Even with a creative research design and a wealth of detailed data, our results should be interpreted in the context of these caveats.

2. Data

We use two data sets to study the influence of HFT on execution costs. The first data set comes from the FSA and identifies HFT activity in the U.K. equity market.

The HFT data are from the FSA Transaction Reporting System (the FSA data set). European legislation, the Markets in Financial Instruments Directive (MiFID), and Chapter 17 of the FSA Handbook define the reportable securities and authorized firms have to report transactions on those securities to the FSA.[7] We focus on the

[6] In August 2010 coverage falls to 40% as some HFTs become direct members of a trading venue and are no longer obliged to report.

[7] The Transaction Report User Pack gives full details of the content: http://www.fsa.gov.uk/pubs/other/trup.pdf

equities market. Only entities subject to FSA regulation must report, though the organization also receives transaction reports from other EEA regulators. Given current regulation, not all HFTs are required to file transaction reports.[8]

The FSA data set provides many variables of interest. It includes the date and time stamp (to the second) of when a trade occurs, the number of shares traded, the counterparty, whether it is a buy or sell trade, and the price at which the trade occurred.[9] Importantly, it includes the user identification (at the firm level) carrying out the trade.[10] As reporting mistakes occasionally occur we winsorize the reported traded volumes in the FSA data at the top 5% to remove extreme values that are likely erroneous. The FSA data set provides an accurate measure of HFT activity from EEA-authorized firms in the U.K. equities asset class. The FSA data set is from November 5, 2007 to August 5, 2011.

HFTs mainly trade in the most liquid stocks, and so we restrict our analysis to the 250 stocks with the largest market capitalization as of November 1, 2007, from the FTSE (the FTSE 100 stocks plus the 150 stocks from the FTSE 250 with the highest market capitalization). For methodological reasons explained in Section 3, we group these stocks into seven groups. The seven categories are based on the market capitalization of the stocks as of November 1, 2007, from Bloomberg and are as follows:

Stock size category	Market capitalization (1 = largest)
1	1–10
2	11–30
3	31–50
4	51–100
5	101–150
6	151–200
7	201–250

Finally, total daily trading volume and stock market capitalization come from Bloomberg.

[8] "... not all high frequency traders are currently required to be authorised under MiFID as the exemption in Article 2.1(d) of the framework directive for persons who are only dealing on own account can be used by such traders." http://ec.europa.eu/internal_market/consultations/docs/2010/mifid/consultation_paper_en.pdf

[9] While the FSA data set has time stamps, we choose not to use them due to questions about the accuracy of the time data and instead focus on day level analysis.

[10] Each report also gives the name of the instrument, who conducted the transaction (the reporting firm), with whom (counterparty 1), and in the case of an agency trade, on behalf of whom (counterparty 2). It discloses the name of the trading platform on which the transaction was made or whether it was off-exchange. To calculate our measure of HFT activity we consider all the transaction reports where an HFT firm reports a principal transaction or is reported as counterparty 1 or counterparty 2. For more information, see the working paper version of this paper (Brogaard, Hendershott, Hunt, Latza, Pedace and Ysusi, 2012).

The second data set is from Abel Noser and documents the execution costs of institutional investors. The Abel Noser data set has been used in several other academic papers.[11] Anand, Irvine, Puckett and Venkataraman (2012) provide a thorough description of the U.S. equities data set in their appendix. Alleviating a potential concern of using the data, Anand, Irvine, Puckett and Venkataraman (2012) argue that survivorship bias is not an issue for two reasons. First, they were reassured by Abel Noser representatives that there was no such bias. Second, they observe firms in the data that dropped out of the sample in the middle of the data set time series. They also show that institutions in the Abel Noser database are representative of 13F institutions.

In the United Kingdom, FTSE 250 during the second quarter of 2010 Abel Noser captures 3.6% of trading volume from 115 unique institutions. For an average stock-day we capture 73,611,604 pounds of trading through the Abel Noser data set. This represents trading of 5.897 million shares for a typical stock-day. On a typical stock-day there are 21 Able Nobel observations with an average trade size of 3.5 million pounds. The median execution shortfall for a trade is 11.3 basis points. There is wide variation with the lower 25th quartile having a cost of only 0.4 basis points and the 75th quartile having a cost of 27.8 basis points.

Abel Noser collects information about institutional trading costs. Its data set contains the date and time of trades by institutions that report to Abel Noser. For each institutional order, the trade price, number of shares, and the direction of the trade are reported. The data set also includes benchmark price measures, such as the value-weighted asset price over the previous trading day and the current day, the end of day price, and the beginning of day price.

We use the market-adjusted execution shortfall of daily institutional traders as our measure of execution costs.[12] The measure can be interpreted as the volume-weighted average price institutional investors pay for a share compared to its true price, the price that prevailed in the market when the sell-side broker received the order.[13] The daily institutional traders' cost of trading for each stock is TC_{jt},

$$TC_{jt} = \sum_{n=1}^{N} \omega_{jtn} \left[buy_{jtn} \left(\frac{P_{jtn} - P_{j,t-}}{P_{j,t-}} \right) - R_{t,FTSE} \right], \tag{1}$$

[11] Some of these papers include: Anand, Irvine, Puckett and Venkataraman (2012, 2013), Chemmanur, He and Hu (2009), Goldstein, Irvine, Kandel and Wiener (2009), and Hu (2009).

[12] Using nonmarket adjusted implementation shortfall gives similar results.

[13] The price that prevailed in the market when the sell-side broker received the order is a variable included in the Abel Noser data set, not one that we create based on intraday time stamps. The Abel Noser data set does include time stamps. However, there is evidence they are not precise.

where n identifies a specific share traded, buy_{jtn} takes the value one if on day t, for stock j, share n was bought by the institutional investor, and negative one if the institutional investor sold share n; P_{jtn} is the price at which the share n for stock j was traded on day t; and $P_{j,t-}$ is the price of stock j at the time the broker received the order; ω_{jtn} is the volume weight. Following Keim and Madhavan (1995) we control for market movements by subtracting the daily return on the FTSE 100 index, $R_{t,FTSE}$, from an order's execution cost after accounting for an order's direction.

Our measure of execution cost focuses solely on market costs (capturing the bid–ask spread, market impact and price drift while executing the order) and does not attempt to account for other trade-related costs, such as brokerage commissions. Note that the execution cost measure can be negative. Negative execution costs have previously been documented in the literature (Keim, 1999). Negative execution costs occur when an institution desires to buy (sell) shares of stock j and j's stock price decreases (increases) between the time the institution gives its broker the order and the time the broker carries out the trade, the transaction would be recorded as having a negative execution cost. An institution that uses limit orders or follows a contrarian strategy should have negative trading costs. While the *TC* measure is on average positive, it does occasionally take on negative values.

2.1. *Execution cost summary statistics*

Table 1, Panel A provides summary statistics of our average daily execution cost measure for the 250 stocks and for each of the seven categories (based on size) in basis points. The daily execution cost is estimated for each stock on each trading day. The reported statistic is the equally weighted average for all stock-days in each category. The first row reports the overall average, while the subsequent seven rows report the average execution cost within the seven size categories. Column 1 reports the mean execution cost, column 2 the median, and column 3 the standard deviation and column 4 the number of observations. Note that the standard deviation of the daily series is high relative to the mean for all the groups. The variation is large enough that we are unable to conclude statistical significance between the means of the different groups.

While Table 1, Panel A documents the cross-sectional variation in execution costs, Figure 1 reports the time series variation.

Figure 1 shows the quarterly average of the execution cost for the FTSE top 250 together with its one-year moving average trend. The quarterly average execution cost is calculated as the equally weighted stock-day execution cost of institutional traders each day in the quarter. The graph on the left shows all quarters, the graph on the right removes observations during the financial crisis.

Figure 1 shows that the execution costs for institutional investors in U.K. equities have a decreasing long-term trend between 2003 and 2011. This downward trend was

Table 1

Summary statistics

The table provides summary statistics. Panel A reports the average daily institutional traders' execution cost for the FTSE top 250 and each of the seven categories (in basis points). FTSE 1–10 represents the top 10 largest stocks. The daily execution cost is estimated for each stock on each trading day. The reported statistic is the equally weighted average for all stock-days in each category. The daily measure of execution cost is the market-adjusted execution shortfall of daily institutional traders: The execution cost is the daily average of the cost of trading for each stock, defined as: $TC_{jt} = \sum_{n=1}^{N} \omega_{jtn}[buy_{jtn}(\frac{P_{jtn}-P_{j,t-}}{P_{j,t-}}) - R_{t,FTSE}]$, where n identifies a specific share traded, buy_{jtn} takes the value one if on day t, for stock j, share n was bought by the institutional investor, and negative one if the institutional investor sold share n; P_{jtn} is the price at which the share n for stock j was traded on day t; and $P_{j,t-}$ is the price of stock j at the time the broker received the order; ω_{jtn} is the volume weight. Panel B characterizes the trading of the 52 identified high-frequency tradings (HFTs). The first row reports the distribution of the number of stocks HFTs trade in. The second row shows the distribution of the daily HFTs volume per stock. The daily volume per stock is calculated as the equally weighted stock-day volume for each HFT. The last row displays the distribution of the daily HFT overall volume which is calculated as the equally weighted daily total volume for each HFT.

Panel A: Market-adjusted execution shortfall

	Mean	Median (in basis points)	Standard deviation	Number of observations (stock-day)
FTSE 1–250	15	14	28	86,652
FTSE 1–10	12	11	60	5,758
FTSE 11–30	12	11	43	13,976
FTSE 31–50	13	11	47	10,315
FTSE 51–100	15	13	41	20,070
FTSE 101–150	20	18	45	17,228
FTSE 151–200	20	17	55	11,929
FTSE 201–250	19	15	74	7,376

Panel B: HFT firms

	Mean	25th %	Median	75th %	Standard deviation
Number of stocks	102	18	95	187	84.4
HFTs' daily volume per stock (shares)	66,236	2,564	9,979	35,939	176,176
HFTs' daily overall volume (millions of shares)	9.129	0.019	0.151	2.365	33.505

temporally interrupted during the financial crisis with costs increasing. Anand, Irvine, Puckett and Venkataraman (2013) find a similar downward trend over time, with an increase in execution fall during the financial crisis. Visually, it is easier to identify the downward trend when excluding the financial crisis (Fig. 1, right).[14,15]

[14] We excluded a two-year period, from July 2007 to June 2009.

[15] The downward time trend is significant at the 5% level using a simple linear regression.

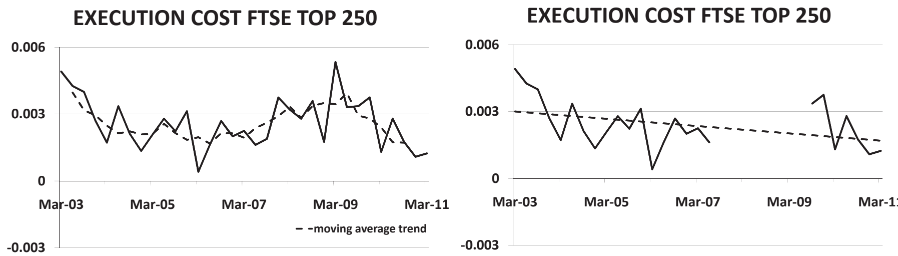

Figure 1

Execution costs over time

The left plot shows the quarterly average of the execution cost for the FTSE top 250 together with its one-year moving average trend. In the right plot we excluded the financial crisis to better identify the execution costs' downward trend. The daily measure of execution cost is the market-adjusted execution shortfall of daily institutional traders (Equation [1]).

2.2. *HFT activity summary statistics*

As HFT lacks a precise common definition, we define which firms primarily engage in HFT based on a number of criteria that appears common in the literature. The primary criterion is based on a definition that HFTs are a subset of algorithmic trading participants that use proprietary capital to generate returns using computer algorithms and low latency infrastructure. Using these criteria, the FSA, along with the three major trading venues for U.K. stock—LSE, BATS and Chi-X—agreed on a set of market participants that were HFT firms. The decision was based on the platforms' understanding of the business of the participant. Fifty-two participants were classified as HFTs. See Alampieski and Lepone (2013) for more details.

In Table 1, Panel B we report summary statistics that characterize the 52 HFTs. Row 1 reports the distribution of the number of stocks HFTs trade in. On average, the HFTs in our sample trade in around 100 stocks out of the 250 stocks considered. Row 2 shows the distribution of the daily HFTs volume per stock. The daily volume per stock is calculated as the equally weighted stock-day volume for each HFT. The last row gives the distribution of the daily HFT overall volume which is calculated as the equally weighted daily total volume for each HFT.

The distribution of all the HFT activity measures are skewed to the right. There are a significant number of HFTs that trade small daily volumes. This may either be due to data limitations (as discussed earlier) or it may be that HFT activity is highly concentrated. In the end, we use aggregate HFT activity, not the firm-by-firm data, and so do not attempt to explain the cross-sectional variation among HFTs.

In untabulated results, we repeat the analysis in this paper based on a list of firms that the supervision division of the FSA has identified as firms that engage in HFT. To this list, we add firms that the authors know trade frequently, rapidly, and tend to have very tight inventory controls. We also identify HFTs in the FSA data set based on observed trading patterns, and analyze the reported trading activities of these traders. This second list consists of 25 firms.[16] The results from this alternative list of firms are qualitatively similar to those using the main list.

Neither approach captures all HFTs. For instance, if a firm engages in multiple trading activities and HFT is not its primary function (such as a large investment bank), we do not consider the firm to be an HFT firm. Second, we miss HFT activity coming from nonauthorized firms that are not subject to EEA regulation based on MiFID and do not trade through a broker.

To corroborate the extent of the coverage of HFT activity in the FSA data set, we compare the FSA data set level of HFT activity to the level of HFT activity in a data set provided to the FSA by the LSE, BATS, and Chi-X—the three primary trading venues for U.K. equities—for a select number of trading days in 2010.[17,18]

At the beginning of 2010 the FSA data cover between 70% and 80% of HFT trading volume, in August 2010 there is a drop in coverage, and by the end of 2010 the FSA data include only 40% of HFT volume. The fall occurs primarily because we do not observe unregulated HFTs that are direct members of trading venues and some HFTs became direct members at this time.[19]

Using the FSA data set, we measure HFT activity (H_{jt}) by its fraction of overall daily volume (from Bloomberg) for each stock on each day,

$$H_{jt} = \frac{HFT\ Vol_{jt}}{Vol_{jt}}, \tag{2}$$

where $HFT\ Vol_{jt}$ is the daily volume traded by the HFTs in stock j on day t and Vol_{jt} is twice the total daily volume traded (once for the buyer and once for the seller) in that same stock j on day t. Note, if an HFT is trading with another HFT then *HFT Vol* will be twice as large as when an HFT is trading with a non-HFT.[20]

[16] Seventeen firms are included in both classification procedures.

[17] The data set used to check our classification could not be used for this study as it is only for a randomly selected number of days in 2010. Attempts to expand the data set were unsuccessful.

[18] The trading venues' data set gives all trades executed from 11:30 to 15:30. In the FSA data, firms must strive to capture the trading time correctly but there are several reasons why the reported time may not be the traded time. Therefore, the volume captured in the FSA data set compared to the trading venues' data may be underestimated. For example, if the trading time is unknown, the default time in the FSA data set is 00:01:00. Or if an external broker fills one order in several transactions, the time reflects the time when the firm becomes the beneficial owner. Misreporting in the FSA's data set can also be a cause of discrepancy between both data sets. The comparison will be affected if the firm code, the counterparty code, the instrument, the venue, the transaction time or the quantity is misreported.

[19] We obtain similar qualitative results when repeating the analysis using only the HFT firms that do not drop out of the data set in August 2010.

[20] Another measure of HFT daily volume participation is the volume of transactions (or shares traded) where an HFT firm is on at least one side of the trade divided by the daily total volume. This measure will

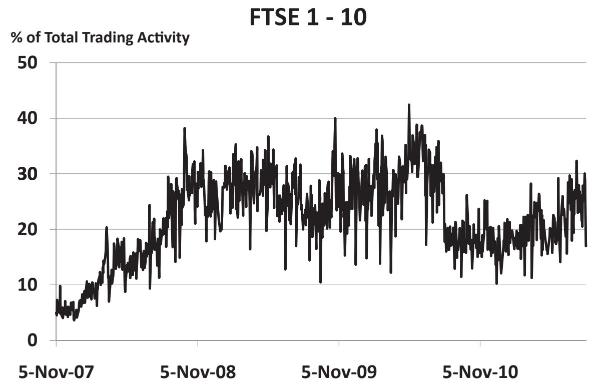

Figure 2

HFT over time in the FTSE top 10

The figure shows the average fraction of trading volume from HFTs in the 10 largest (by market capitalization) stocks in the United Kingdom between November 2007 and August 2011. The fraction is defined as $H_{jt} = HFT\ Vol_{jt}\ /\ Vol_{jt}$, where $HFT\ Vol_{jt}$ is the daily volume traded by HFTs in stock j on day t and Vol_{jt} is twice the total daily volume traded (once for the buyer and once for the seller) in that same stock j on day t.

Figure 2 graphs the average HFT activity, H_{jt} for the largest 10 U.K. equities for the entire FSA data set.

Figure 2 shows that HFT activity increases steadily from the beginning of our sample, November 2007, until early 2009. From early 2009 until August 2010 HFT activity fluctuates around 30% of trading volume. In August 2010 HFT activity drops to about 20% of trading volume. The decline is driven by HFTs that changed from sponsored access through regulated brokers to direct access to the trading venues. Afterward we see a constant level of participation with a slight increase at the end of the sample.[21] The HFT activity we capture is in line with Tabb Group, a data-gathering agency that claims that European HFT has increased from 5% in 2006 to 20% in 2008, to almost 40% in 2011.[22]

Table 2 breaks down by market capitalization the amount of trading activity coming from HFT.

---

yield higher estimates than the one used in this paper. If, for example, there were no transactions where HFTs were on both sides of the trade, this measure will be twice our measure.

[21] We find a similar pattern for the largest 100 stocks. For stocks in the categories 101–250 HFT activity is more stable over time.

[22] http://www.ft.com/cms/s/0/74ace24a-ac00–11e0-b85c-00144feabdc0.html#axzz1oX6Spjlj

Table 2

Cross-section of high-frequency trading (HFT) activity

The table gives summary statistics of our measure of HFT activity, daily volume participation, for the FTSE top 250 largest stocks and each of the seven categories. The seven categories are based on stock market capitalization. FTSE 1–10 represents the top 10 largest stocks. HFT activity is defined as $H_{jt} = HFT\ Vol_{jt}/Vol_{jt}$, where $HFT\ Vol_{jt}$ is the daily volume traded by the HFTs in stock j on day t and Vol_{jt} is twice the total daily volume traded (once for the buyer and once for the seller) in that same stock j on day t.

	Mean (%)	Median (%)	Standard deviation	Largest observation (%)
FTSE 1–250	11.25	11.40	4.05	22.90
FTSE 1–10	22.09	17.06	7.56	42.42
FTSE 11–30	14.55	11.18	4.87	26.48
FTSE 31–50	12.83	8.54	5.29	25.98
FTSE 51–100	10.63	7.32	4.27	20.11
FTSE 101–150	8.15	5.98	3.63	19.33
FTSE 151–200	5.70	4.16	2.59	17.12
FTSE 201–250	4.45	2.60	2.65	15.58

Table 2 reports the stock-day average HFT activity, H_{jt} based on the size category. Column 1 reports the mean, column 2 the median, column 3 the standard deviation, and column 4 the largest HFT activity, H_{jt}, for a stock-day in the associated category. The first row shows the overall participation rate for all 250 stocks. The next seven rows report the HFT activity for the different size categories.

As the size of the stock decreases, the level of HFT also decreases, and the relationship is monotonic. HFTs are involved in 22.09% of trades for the top 10 stocks, but only 4.45% for the smallest 50 stocks.

3. Empirical methodology and results

In this section, we examine the impact of technological upgrades in the LSE on HFT activity and the execution cost of institutional investors separately. We present figures of HFT trading and execution costs around the technology change events and we also run panel regressions to examine the effect. Finally, to assess the link between HFT and institutional investors' execution costs we implement a 2SLS regression.

3.1. The technological change: Exchange latency changes

The technological upgrades we base our analysis on are latency changes by the LSE. Given that the changes in network speeds are in milliseconds, the changes will only have a direct impact on computer-based traders. Non-HFT algorithmic traders may be marginally impacted by millisecond latency changes, but arguably those who depend most on speed, HFTs, are the most likely to be affected. In addition, while HFTs may lobby the exchange to decrease its latency, HFTs do not determine exactly when such changes are implemented. As a result, network latency may provide a reasonable shock to HFT activity while having little direct impact on the trading activity of institutional investors. Wagener, Kundisch, Riordan, Rabhi, Herrmann and Weinhardt (2010) follow a similar approach.

Since 2007 there have been a variety of technology changes at the LSE reducing latency. From the 2011 Annual Report of the LSE, we collect a list of five technology upgrades during the sample that decrease the latency of the fastest traders from 11 ms to 0.113 ms. The upgrades include the major changes of the TradElect system and the introduction of the Millennium system:

System	Implementation date	Latency (milliseconds)	Percent decline
TradElect 2	October 31, 2007 (before the sample)	11	
TradElect 3	September 1, 2008	6	45%
TradElect 4	May 2, 2009	5	17%
TradElect 4.1	July 20, 2009	3.7	26%
TradElect 5	March 20, 2010	3	19%
Millennium	February 14, 2011	0.113	96%

We study the implementations of TradElect 3, 4, 4.1, and 5. Originally, the research design called for the use of all five latency improvements during the sample; however, due to data limitations and market conditions, we are limited to the TradElect 3 to 5 upgrades. The TradElect 2 implementation occurs prior to our data set and gives us the baseline latency at the onset of our sample. The Millennium implementation occurs during a time in which the fraction of HFT activity captured in the FSA data set has declined, and the measure may not accurately capture changes in HFT activity.

3.2. The impact of the technology change on HFT activity and execution costs

To analyze the impact of technology changes on HFT activity and long-term investors' execution costs, we examine trading around the technology change events. We graphically compare the levels of the daily average of our HFT activity measure for the FTSE top 250 before and after the latency upgrade implementation change. We plot the time series of the stock-day average HFT activity, H_{jt}, for a 20-day window, 10 days before and 10 days after the latency change. We repeat the graph but this time measuring the stock-day average execution cost, TC_{jt}. Figure 3 shows HFT activity and execution cost around the TradElect 3 system upgrade.[23]

We use a narrow time window to isolate the possible effect of latency changes from other effects due to prevalent market conditions. While both the HFT activity and

[23] Figures of HFT activity and execution costs around the other system upgrades are available on request from the authors.

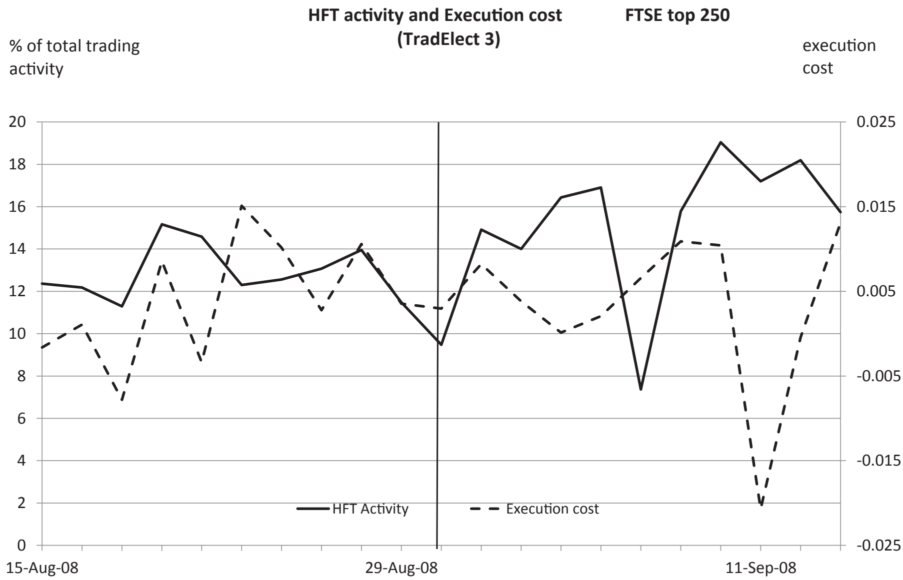

Figure 3

HFT activity and execution costs

The figure shows the average daily level of HFT activity and institutional execution costs among the 250 stocks with the largest market capitalization around TradElect 3. The left axis measures HFT activity. The right axis measures execution costs. The fraction of HFT activity is defined as $H_{jt} = HFT\ Vol_{jt}/Vol_{jt}$, where $HFT\ Vol_{jt}$ is the daily volume traded by the HFTs in stock j on day t and Vol_{jt} is twice the total daily volume traded (once for the buyer and once for the seller) in that same stock j on day t. The execution cost is the daily average of the cost of trading for each stock, defined as: $TC_{jt} = \sum_{n=1}^{N} \omega_{jtn}[buy_{jtn}(\frac{P_{jtn}-P_{j,t-}}{P_{j,t-}}) - R_{t,FTSE}]$, where n identifies a specific share traded, takes the value one if on day t, for stock j, share n was bought by the institutional investor, and negative one if the institutional investor sold share n; P_{jtn} is the price at which the share n for stock j was traded on day t; and $P_{j,t-}$ is the price of stock j at the time the broker received the order; ω_{jtn} is the volume weight.

execution costs graphs are noisy there is a small visual change around the structural break. After the technology upgrade the level of HFT activity appears to increase while the institutional execution cost appears to decrease, although the two do not appear to move in opposite directions on a day-by-day basis.

We formally test whether HFT activity changed after the LSE technology upgrades. We perform the following ordinary least squares panel regression for a 20-day centered window around the TradElect 3, 4, 4.1, and 5 implementations:[24]

[24] We run several robustness checks to ensure that our results are not an artifact of our particular variable and model specifications. Similar results are obtained using a 40-day window. The results are robust to

$$H_{jt} = \alpha_j + \sum_{i=1}^{7} \left(\beta_1^i t + \beta_2^i L_t\right) + \nu V_{jt} + \varepsilon_{jt}, \tag{3}$$

where H_{jt} is our measure of activity by HFTs for stock j on day t, V_{jt} is a control variable: log total volume. We include stock fixed effects (α_j). Standard errors are double clustered by stock and day.[25]

Technological advances may have a varying impact for different types of equities. We capture this potential variation by groupings stocks into seven groups (denoted by superscript i) according to their market capitalization and include seven time and latency coefficients, one for each group. L_t takes the value of the LSE system latency at time t for observations after an LSE trading system, that is, it takes the value shown in the "Latency" column of the table showing system changes and latency above, until the next system change reduces latency: a decreasing step function. Finally, we allow a linear time trend for each group by including t, which takes the value one on the first trading day in the data set and increases incrementally by one for each subsequent trading day.

The use of seven groups reduces the amount of variables to be estimated, which is necessary for the individual regressions using short windows around the TradElect implementations. The technique still allows us to capture the cross-sectional variation in the effect of latency changes. We expect highly liquid stocks to be more affected by technology changes as HFTs tend to be more active in these stocks.

The results are reported in Table 3.[26]

The regression is executed for the TradeElect updates 3, 4, 4.1, and 5. The coefficients (z-statistic) of each regression are reported in columns 1 (2), 3, (4), 5 (6), and 7 (8), respectively. The first seven rows represent the different L_t^i variables. The latency variables used imply that, if the change in latency is 1 ms as with TradElect 4, a coefficient of –0.01 translates to an increase in HFT activity that is one percentage point of traded volume.

While the coefficients on the *LSE latency* variables are generally statistically insignificant for TradElect 3 and TradElect 4.1, most statistically differ from zero in TradElect 4 and TradElect 5. For TradElect 4 all coefficients are negative; three

---

different specifications of latency including log latency and 1/latency. Excluding volume or the linear trend from the regression specification reduces the R^2 and significance, but has no impact on signs of our variables of interest. Additionally, we run placebo tests using four randomly chosen dates away from our current dates and none show equivalent results to what we find for the four actual changes. The results are available on request from the authors.

[25] While both the Abel Noser data set and the FSA data set have time stamps, they are insufficiently accurate for matching the data at the trade level, so we cannot conduct analysis at this level. We tried to line up the Abel Noser and FSA data sets with tick by tick trade data from Thompson Reuters' Sirca database, but were unsuccessful in the exercise.

[26] The estimates for the time trends are not shown in Table 3, but are available on request from the authors. As in all regressions where applicable, the intercept is the average of the stock fixed effects.

Table 3

Latency and high-frequency trading (HFT) volume

The table shows the results from an ordinary least squares panel regression for a 20-day centered window around the TradElect 3, 4, 4.1, and 5 implementations: $H_{jt} = \alpha_j + \sum_{i=1}^{7} (\beta_1^i t + \beta_2^i L_t) + \nu V_{jt} + \varepsilon_{jt}; H_{jt}$ is the measure of activity by HFTs for stock j on day t, V_{jt} is a control variable: log volume. We include stock fixed effects, α_j. We group stocks into seven categories (denoted by superscript i) according to their market capitalization and include seven time and latency variables. L_t takes the value of the LSE system latency at time t for observations after an LSE trading system. We allow a linear time trend for each group by including t, which takes the value one on the first trading day in the data set and increases incrementally by one for each subsequent trading day. Standard errors are double clustered by stock and day.

	TradElect 3		TradElect 4		TradElect 4.1		TradElect 5	
	10 days before and after September 1, 2008		10 days before and after May 2, 2009		10 days before and after July 20, 2009		10 days before and after March 20, 2010	
HFT fraction	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat
FTSE 1–10	0.055	1.39	−0.060	−1.31	−0.002	−0.03	−0.071***	−2.82
FTSE 11–30	0.021	0.67	−0.047	−1.52	0.009	0.28	−0.040***	4.14
FTSE 31–50	−0.013	−0.41	−0.038**	−2.00	0.010	0.36	−0.026**	2.44
FTSE 51–100	0.012	0.58	-0.027	−1.46	0.031	0.73	−0.013*	−1.71
FTSE 101–151	0.024**	2.48	−0.025**	−2.13	0.025	0.91	0.001	0.06
FTSE 151–200	0.006	0.56	−0.015	−1.45	0.024	1.14	0.004	0.90
FTSE 201–250	0.010**	2.21	−0.021***	−2.62	0.033*	1.87	0.008**	2.06
Total volume	−0.010	−0.84	−0.034***	−7.38	−0.032***	−9.83	−0.029***	−8.77
Intercept	0.457	1.28	1.703***	7.08	1.361***	6.64	1.499***	12.92
Adj-*R* squared		0.767		0.809		0.723		0.850
N		4,700		4,546		4,540		4,432

***, **, * indicate statistical significance at the 0.01, 0.05 and 0.10 level, respectively.

groups are significantly negative at the 5% level. For TradElect 5 the three groups covering the largest stocks are significantly smaller than zero (at the 5% level). This indicates that HFT activity increases after a technology upgrade on the LSE. After the 1 ms improvement in minimum latency by TradElect 4, HFT activity jumps up by two to four percentage points.

The 0.7-ms improvement of TradElect 5 increases the share of HFT activity by two to seven percentage points. While the level of the change was the smallest for TradElect 5, in percentage terms it was a 19% decrease in latency, higher than TradElect 4, but less than TradElect 3 and 4.1. Other than speed and HFT being positively associated we do not have a strong prior on the functional form of their relationship. Note that the control variable, *Total Volume*, has a negative coefficient. A larger total trading volume decreases the share of HFT participation. For instance,

Table 4

Latency and execution costs

The table shows the results from an ordinary least squares panel regression for a 20-day centered window around the TradElect 3, 4, 4.1, and 5 implementations: $TC_{jt} = \alpha_j + \sum_{i=1}^{7} (\beta_1^i t + \beta_2^i L_t) + \nu V_{jt} + \varepsilon_{jt}$; TC_{jt} is the institutional investors' execution cost for stock j on day t, V_{jt} is a control variable: log volume. We include stock fixed effects, α_j. We group stocks into seven categories (denoted by superscript i) according to their market capitalization and include seven time and latency variables. L_t takes the value of the London Stock Exchange (LSE) system latency at time t for observations after an LSE trading system. We allow a linear time trend for each group by including t, which takes the value one on the first trading day in the data set and increases incrementally by one for each subsequent trading day. Standard errors are double clustered by stock and day.

	TradElect 3		TradElect 4		TradElect 4.1		TradElect 5	
	10 days before and after September 1, 2008		10 days before and after May 2, 2009		10 days before and after July 20, 2009		10 days before and after March 20, 2010	
HFT fraction	L coef.	z-stat	L coef.	z-stat	L coef.	z-stat	L coef.	z-stat
FTSE 1–10	−0.0050	−0.87	−0.0033	−1.12	0.0005	0.10	0.00003	0.02
FTSE 11–30	0.0017	0.32	0.0053**	2.29	−0.0080*	−1.74	−0.0009	−0.57
FTSE 31–50	0.0029	0.88	−0.0006	−0.24	−0.0099**	−1.97	0.0029*	1.87
FTSE 51–100	0.0068	1.43	−0.0026	−0.98	−0.0052**	−1.96	0.0016	1.21
FTSE 101–151	0.0054	1.16	−0.0027	−0.72	0.0029	0.70	−0.0009	−0.31
FTSE 151–200	−0.0042	−0.82	−0.0009	−0.22	0.0069	1.17	0.0014	0.39
FTSE 201–250	−0.0001	−0.03	−0.0089	−1.59	−0.0026	−0.25	0.0046	0.66
Total volume	0.0011	0.89	0.0010	0.91	0.0008	0.59	0.0009	0.86
Intercept	0.0235	0.66	−0.0036	−0.12	−0.0162	−0.54	−0.0359	−1.23
Adj-R squared		0.104		0.126		0.110		0.231
N		2,189		2,433		2,362		1,159

***, **, * indicate statistical significance at the 0.01, 0.05 and 0.10 level, respectively.

the coefficient for TradElect 3 implies that a 1% increase in trading volume relates to a decrease in HFT fraction of trading by 1%.

We repeat the analysis but with the dependent variable being TC_{jt}, the execution cost measure from Equation (1). The following regression is identical to that in Equation (3) except for the dependent variable:

$$TC_{jt} = \alpha_j + \sum_{i=1}^{7} \left(\beta_1^i t + \beta_2^i L_t\right) + \nu V_{jt} + \varepsilon_{jt}, \tag{4}$$

where TC_{jt} is the institutional investors' execution cost for stock j on day t, and the rest of the regression is as described for Equation (3).

The results are reported in Table 4.

The table reports its results in the same layout as Table 3. Regardless of the stock size category or the system upgrade there is no consistent link between execution cost and latency. Table 4 confirms that there is no clear measurable association between the technology changes and execution costs for the 20-day sample windows.

Table 3 shows that the exchange upgrades resulted in more HFT activity. At the same time, Table 4 shows the size of execution costs stayed the same. However, from these two analyses one cannot conclude that HFT activity does not influence execution costs.

Our aim is to assess the marginal impact a small change (increase) of HFT has on institutional investors' execution costs. We cannot draw conclusions about the causal impact of HFT by simply looking at the association between HFT activity and execution costs. First, some third factor could drive both HFT activity and execution costs. For example, during the period in question, fundamentals, or aspects of the financial crisis, could cause greater uncertainty that both causes HFTs to trade more and causes spreads and execution costs to increase. Second, a correlation between HFT activity and execution costs could be because execution costs affect HFT participation and not the other way around.

3.3. The 2SLS approach

To overcome the endogeneity problem we implement a 2SLS methodology. In the first stage, HFT activity is regressed on the instrument and some control variables to find the relationship between the instrument and HFT activity. The instrumental variable used to isolate this element of HFT activity is the speed change in the LSE's matching engine. In the second stage, we use the outcomes of the first stage to isolate a component of HFT activity that is independent of execution costs.

2SLS requires a variable that satisfies two conditions. First, the variable is correlated with HFT activity. Second, the variable must satisfy the exclusion restriction: it must not be correlated with execution costs except through its relationship to HFT activity. Changes in the LSE's latency could affect algorithmic trading other than HFT. Our identification relies on the assumptions that the few millisecond changes are small enough that only the very fastest traders would be directly affected and that those traders are HFTs.

We regress execution costs on the predicted component of HFT activity (from the first regression) and control variables. The coefficient on HFT activity in this second stage is the causal effect of HFT activity on execution costs. If the exclusion restriction holds, this coefficient is an unbiased estimator for this causal effect.

To conduct the 2SLS regression, in the first stage we regress the level of HFT activity at the stock-day level on a variable capturing the speed change in the LSE's systems to handle electronic messages and relevant control variables, as defined in

Table 5

First stage

The table shows the results from a panel regression with stock fixed effects of the fraction of high-frequency trading (HFT) volume, defined in Equation (2), on (i) London Stock Exchange (LSE) latency, (ii) total volume traded, (iii) a short sale ban dummy, and (iv) linear time trends for each category. The model is estimated using ordinary least squares for 20-day windows around the four TradElect upgrades separately. Standard errors are double clustered at the stock and day level.

	TradElect 3		TradElect 4		TradElect 4.1		TradElect 5	
	10 days before and after September 1, 2008		10 days before and after May 2, 2009		10 days before and after July 20, 2009		10 days before and after March 20, 2010	
HFT fraction	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat
FTSE 1–10	0.0314	1.36	−0.0268	−1.07	−0.0137	−0.32	−0.0554***	−5.00
FTSE 11–30	0.0114	0.40	−0.0537*	−1.78	0.0101	0.30	−0.0368***	−3.60
FTSE 31–50	−0.0185	−0.63	−0.0422**	−2.03	0.0076	0.25	−0.0362***	−2.77
FTSE 51–100	0.0078	0.41	−0.0338	−1.47	0.0374	0.84	−0.0146*	−1.64
FTSE 101–151	0.0259***	2.93	−0.0306**	−2.44	0.0296	1.13	−0.0115	−1.45
FTSE 151–200	0.0024	0.39	−0.0142	−1.36	0.0481	1.32	0.0054	0.61
FTSE 201–250	0.0109*	1.94	−0.0266**	−2.15	0.0840**	2.57	0.0164	1.41
Total volume	−0.0080	−0.46	−0.0399***	−6.20	−0.0415***	−9.15	−0.0419***	−8.66
Intercept	0.0653	0.30	1.1519***	7.25	1.0406***	5.10	0.6029***	7.03
Adj-*R* squared		0.771		0.721		0.626		0.838
N		2,184		2,433		2,362		1,159

***, **, * indicate statistical significance at the 0.01, 0.05 and 0.10 level, respectively.

Equation (3). The first stage regression includes a linear time trend. The model is estimated for 20-day windows around the four TradElect upgrades separately.

In the second stage, we use the estimated proxy for HFT activity from the first stage in a regression with the dependent variable being the stock-day execution costs of a set of firms included in the Abel Noser data set:

$$TC_{jt} = \alpha_j + \sum_{i=1}^{7} \beta_1^i t + \theta \hat{H}_{jt} + \nu V_{jt} + \varepsilon_{jt}, \tag{5}$$

where TC_{jt} is the institutional investors' execution cost for stock j on day t, and $\hat{H}_{jt}$ is the predicted measure of HFT activity from the first stage regression. We use the same control variable as in the first stage regression (log trading volume (V_{jt})), allow for fixed effects at the stock level (α_j), and a linear time trend. Again seven groups (i) are used to capture differences in the cross-section.

Table 5 summarizes the results from the first stage regressions. With few exceptions across the group sizes and exchange upgrades, the latency coefficients are either negative and statistically significantly different than zero, or not statistically significant than zero. As the system upgrades reduce the size of the latency variable,

Table 6

Second stage

The table shows the result from the second stage in a two-stage least-squares (2SLS) regression of execution costs on instrumented high-frequency trading (HFT) activity for the TradElect 4 and TradElect 5 events. The following ordinary least squares regression is performed: $TC_{jt} = \alpha_j + \sum_{i=1}^{7} \beta_1^i t + \theta \hat{H}_{jt} + \nu V_{jt} + \varepsilon_{jt}$. TC_{jt} is the institutional investors' execution cost for stock j on day t, and $\hat{H}_{jt}$ is the predicted measure of HFT activity from the first stage regression. We use the same control variable as in the first stage regression (log trading volume (V_{jt})), allow for fixed effects at the stock level (α_j), and a linear time trend. Seven groups (i) are used to capture differences in the cross-section. Standard errors are double clustered at the stock and day level.

Second stage				
	TradElect 4		TradElect 5	
T-cost	L coef.	*z*-stat	*L* coef.	*z*-stat
Predicted HFT	0.0151	0.44	–0.0054	–0.23
Total volume	0.0019	1.37	0.0008	0.63
Intercept	−0.0281	−1.12	−0.0109	−0.49
Adj-*R* squared		0.0018		0.0058
N		2,433		1,159

a negative coefficient means that as latency decreases, HFT increases. For instance, the coefficient on *FTSE 11–30* for TradElect 4 is –0.0537 which says that when the exchange latency decreases by 1 ms, HFT fraction of trading in stocks that are between the 11th and 30th largest in size increases by 5%. For most of the categories in TradElect 4 and TradElect 5 the exchange upgrade has a statistically significant impact on the level of HFT activity. We focus on these two upgrades for the second stage.

The results for the second stage are presented in Table 6.

The variable of interest is the coefficient on *Predicted HFT*. In both of the system upgrade events we find no statistically significant relationship between the instrumented HFT activity variable and execution costs.[27] Neither coefficient is near being statistically significant (*z*-stat of 0.44 and –0.23 for TradElect 4 and TradElect 5, respectively).[28]

The changes being analyzed occurred during the financial crisis, a period in which execution costs were highly volatile. Given our measure of HFT only increases by 2% it may be that the impact is too small to detect. Before concluding that we

[27] The second stage results are omitted for TradElect 3 and 4.1 from Table 6 because of potential confounding events. The TradElect 3 event occurs just as the financial crisis is peaking, and TradElect 4.1 in the 40-day window pre-period analysis overlaps with the TradElect 4 40-day window post-period analysis.

[28] We repeat the analysis using 40-day windows and the results are qualitatively similarly.

cannot reject that HFT does not impact institutional investors' execution costs we repeat the analysis with two significant changes.

4. Robustness

We implement two robustness checks. First, we pool the four event studies into one analysis. Second, we analyze portfolios instead of stocks. Both approaches produce similar findings as our main specification.

4.1. Pooled regression

To increase the power of the event-based regressions, we pool the four events into a single regression. The drawback is we restrict the slope coefficient on the latency changes to be the same. The basic 2SLS regression setup as described in Equations (3) and (5) does not change. However, to limit the number of coefficients that need to be estimated we restrict the slope coefficients on the seven groups to be the same. We include event-window dummy variables,

$$H_{jt} = \alpha_j + w_k + \sum_{i=1}^{7} \beta_2^i L_t + \nu V_{jt} + \varepsilon_{jt}, \tag{6}$$

$$TC_{jt} = \alpha_j + w_k + \theta \hat{H}_{jt} + \nu V_{jt} + \varepsilon_{jt}, \tag{7}$$

where k runs from one to four, w is a window dummy taking the value one if an observation is during the event-window k. The other variables are defined as before. The results are summarized in Table 7.

Table 7, Panel A reports the first stage results and Table 7, Panel B reports the second stage. There are fewer statistically significant coefficients in the first stage in the pooled analysis than in the main specification. Only *FTSE 11–30* is statistically significantly different than zero. Like in the main specification though, most of the coefficients are negative.

In the second stage, the results are inconclusive. We again find no relationship between HFT and institutions' execution costs.

4.2. Portfolio regressions

To reduce noise in the HFT and execution cost measures, the 2SLS regressions are run on seven portfolios instead of 250 stocks. The portfolios are formed by stock market capitalization as defined in Section 3. Within each portfolio the values of HFT activity (H) and execution costs (TC) are weighted by market trading volume. The regression setup is as described in Equations (3) and (5) with the exception

Table 7

Pooled regression

The table shows the result from an alternative two-stage least-squares (2SLS) regression specification. Panel A reports the first stage. Panel B reports the second. The specification is identical to that used in Table 5 for the first stage and Table 6 for the second stage except all observations are pooled into a single regression and event-window dummy variables are included. The first stage regression is: $H_{jt} = \alpha_j + w_k + \sum_{i=1}^{7} \beta_2^i L_t + \nu V_{jt} + \varepsilon_{jt}$. The second stage is: $TC_{jt} = \alpha_j + w_k + \theta \hat{H}_{jt} + \nu V_{jt} + \varepsilon_{jt}$, where k runs from one to four, w is a window dummy taking the value one if an observation is during the event-window k. The other variables are defined as before. Standard errors are double clustered at the stock and day level.

Panel A: First stage			*Panel B: Second stage*		
HFT fraction	*L* coef.	*z*-stat	*T*-cost	*L* coef.	*z*-stat
FTSE 1–10	−0.0116	−1.45	HFT-hat	−0.0203	−0.69
FTSE 11–30	−0.0112*	−1.72	Total volume	0.0004	0.51
FTSE 31–50	−0.0032	−0.46	Intercept	0.0199	1.06
FTSE 51–100	−0.0012	−0.19	Adj-*R* squared		0.0306
FTSE 101–151	−0.0082	−1.3	*N*		8,138
FTSE 151–200	−0.0008	−0.11			
FTSE 201–250	−0.0051	−0.74			
Total volume	−0.0259***	−6.34			
Intercept	0.3416***	4.67			
Adj-*R* squared		0.506			
N		8,138			

***, **, * indicate statistical significance at the 0.01, 0.05 and 0.10 level, respectively.

that variables are now defined on the portfolio level ($i = 1, \ldots 7$) rather than on the individual stock level ($j = 1, \ldots, 250$):

$$H_{it} = \alpha_i + \sum_{i=1}^{7} \left(\beta_1^i t + \beta_2^i L_t \right) + \nu V_{it} + \varepsilon_{it}, \tag{8}$$

$$TC_{it} = \alpha_i + \sum_{i=1}^{7} \beta_1^i t + \theta \hat{H}_{it} + \nu V_{it} + \varepsilon_{it}. \tag{9}$$

The results are summarized in Table 8.

Table 8, Panel A reports the first stage results and Panel B the second stage. The results are qualitatively similar to the main specification. In the first stage only TradElect 4 and TradElect 5 have consistently statistically significant coefficients, and the coefficient signs are normally negative, meaning HFT increased after exchange speed upgrades.

In the second stage, the results are again inconclusive. We again find no relationship between HFT and institutions' execution costs.

Table 8

Portfolio regression

Table 8 shows the result from an alternative two-stage least-squares (2SLS) regression specification. Panel A reports the first stage. Panel B reports the second. The specification is identical to that used in Table 5 for the first stage and Table 6 for the second stage except the unit of observation is a portfolio, based on the seven stock size categories. Within each portfolio, the values of high-frequency trading (HFT) activity (*H*) and execution costs (*TC*) are weighted by market trading volume. Standard errors are double clustered at the stock and day level.

Panel A: First stage

	TradElect 3		TradElect 4		TradElect 4.1		TradElect 5	
HFT fraction	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat	*L* coef.	*z*-stat
FTSE 1–10	0.0423***	3.65	0.0059	0.45	−0.0350	−1.09	−0.0249***	−3.80
FTSE 11–30	0.0013	0.06	−0.0251	−1.53	0.0171	0.54	−0.0466***	−9.90
FTSE 31–50	−0.0413**	−2.08	−0.0252	−1.27	0.0083	0.20	−0.0155*	−1.91
FTSE 51–100	0.0151	1.18	−0.0269*	−1.70	0.0212	0.52	−0.0112*	−1.94
FTSE 101–151	0.0042	0.32	−0.0331**	−2.30	0.0085	0.26	−0.0080*	−1.88
FTSE 151–200	0.0130	1.21	−0.0047	−0.36	0.0215	0.63	−0.0101	−1.30
FTSE 201–250	−0.0028	−0.23	−0.0189	−1.52	0.0396*	1.81	−0.0002	−0.06
Total volume	0.0353*	1.67	−0.0170	−1.10	−0.0398***	−3.30	−0.0432***	−4.35
Intercept	−0.7804*	−1.86	0.5373*	1.69	1.1450***	3.88	1.1202***	5.93
Adj-*R* squared		0.871		0.898		0.828		0.940
N		140		140		140		139

Panel B: Second stage

	TradElect 4		TradElect 5	
T-cost	*L* coef.	*z*-stat	*L* coef.	*z*-stat
HFT-hat	−0.0314	−0.31	0.0526	0.71
Total volume	0.0042	1.64*	0.0021	0.48
Intercept	−0.0757	−1.4	−0.0441	−0.50
Adj-*R* squared		0.0325		0.0075
N		140		139

***, **, * indicate statistical significance at the 0.01, 0.05 and 0.10 level, respectively.

5. Conclusion

If transaction costs are low, market participants are better able to hold the assets most suited to them, and informed participants are more able to trade on their private information and impound it into asset prices. Thus, it is important to understand how new developments in financial market microstructure impact institutional transaction costs.

HFT has quickly become a term known to the general public. The idea of computers running financial markets has raised concerns among other market participants, the media, regulators, academics, and the general public. One of these concerns is that HFT has increased execution costs, a component of transaction costs, at least for some market participants.

We show that in the United Kingdom, like in the United States, there has broadly been a decrease in institutional execution costs over the last decade. This trend, however, was interrupted by the financial crisis, which caused execution costs to increase between mid 2007 and mid 2009. Because HFT increased substantially during the financial crisis some have asserted that HFT may be responsible. We provide evidence on whether HFT causally increases institutional execution costs by studying the major latency changes made by the LSE. We find an association between the latency changes and HFT activity but no measurable association between these latency changes and execution costs. Formally using the latency changes as an instrumental variable shows no link between HFT and execution costs.

Institutional execution costs are lower in 2011 than in the early 2000s while HFT is substantially higher. However, we fail to find evidence that HFT is responsible for the decline in execution costs. It may be that automation in the trading process unrelated to HFT benefits institutions and lowers their costs. It is also possible that our tests are limited because we examine causality in relatively narrow periods around technology changes. Noise in our measure of execution costs could overwhelm possible effects: intraday prices are noisy which makes execution cost measures have high variance. In addition, the changes in the level of HFT we find are relatively small and take place from an already relatively high level of HFT. If it is the introduction of HFT and not a mild increase in HFT that increases execution costs, then our approach would be unable to detect it.

Understanding how the evolution of technology and financial markets influences trading costs is relevant to both regulators and investors. The general decline in execution costs and our failure to find a relationship between HFT and execution costs do not support the need for strong regulation of HFT. At the same time, the inability to show clear benefits of technological advances suggests that further study is warranted. In particular, a more detailed examination of whether some types of institutions are disadvantaged is crucial. If any disadvantaged institutions are especially important, for example, pension funds or long-term investors helping to uncover fundamental information about companies, then a clearer case for market structure reform or monitoring can be made.

References

Alampieski, K. and A. Lepone, 2013. High frequency trading in U.K. equity markets: Evidence surrounding the US market open. *Working paper*, University of Sydney.

Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2012. Performance of institutional trading desks: An analysis of persistence in trading costs, *Review of Financial Studies* 2(25), 557–598.

Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2013. Institutional trading and stock resiliency: Evidence from the 2007–09 financial crisis, *Journal of Financial Economics* 108(3), 773–797.

Benos, E. and S. Sagade, 2012. High-frequency trading behaviour and its impact on market quality: Evidence from the UK equity market. *Working paper*, Bank of England.

Brogaard, J., T. Hendershott, S. Hunt, T. Latza, L. Pedace, and C. Ysusi, 2012. High-frequency trading and the execution costs of institutional investors. U.K. Foresight Project *Working paper*.

Brogaard, J., T. Hendershott, and R. Riordan, 2013. High frequency trading and price discovery, *Review of Financial Studies*, forthcoming.

Carrion, A., 2013. Very fast money: High-frequency trading on the NASDAQ, *Journal of Financial Markets* 16(4), 680–711.

Chemmanur, T., S. He, and G. Hu, 2009. The role of institutional investors in seasoned equity offerings, *Journal of Financial Economics* 94(3), 384–411.

Goldstein, M., P. Irvine, E. Kandel, and Z. Wiener, 2009. Brokerage commissions and institutional trading patterns, *Review of Financial Studies* 22(12), 5175–5212.

Gondat-Larralde, C. and K. James, 2008. IPO pricing and share allocation: The importance of being ignorant, *Journal of Finance* 63(1), 449–478.

Hu, G., 2009. Measures of implicit trading costs and buy–sell asymmetry, *Journal of Financial Markets* 12(3), 418–437.

Keim, D., 1999. An analysis of mutual fund design: The case of investing in small-cap stocks, *Journal of Financial Economics* 51(2), 173–194.

Keim, D. and A. Madhavan, 1995. Anatomy of the trading process: Empirical evidence on the behavior of institutional traders, *Journal of Financial Economics* 37(3), 371–398.

Malinova, K., A. Park, and R. Riordan, 2013. Do retail traders suffer from high frequency traders. *Working paper*, University of Toronto.

Menkveld, A., 2013. High frequency trading and the new market makers, *Journal of Financial Markets* 16(4), 712–740.

Wagener, M., D. Kundisch, R. Riordan, F. Rabhi, P. Herrmann, and C. Weinhardt, 2010. Price efficiency in futures and spot trading: The role of information technology, *Electronic Commerce Research and Applications* 9(5), 400–409.