Fast Aggressive Trading*

Torben Latza, Ian W. Marsh and Richard Payne†

September 12, 2017

Abstract

We subdivide trades on the London Stock Exchange according to their reaction times. We show that faster trades are associated with smaller execution costs than slower trades. However, most fast trades (reaction times less than one second) lead to virtually no permanent price impact while slow trades have the usual positive long run price impacts. The very fastest category of trades - reaction times of less than one millisecond - are associated with large execution costs when compared with other fast trades, but also have a large price impact comparable in magnitude with slow trades. Slow trades and the very fastest trades are therefore comparable across these two key dimensions, while other fast trades are very different. We find no evidence supporting the supposition that counter-parties can be manipulated into trading against faster traders. Most fast traders are fairly innocuous, managing to reduce trading costs without any adverse selction or manipulation concerns. The very fastest trades pay a considerable amount to execute but bring information to the market. Together, these results suggest that regulating trading speed in order to curb issues thought to be associated with high-frequency trading may be at best a

*Thanks to Cass Business School for supporting this work with a research pump-priming grant. Early versions of this work were presented at the WBS Frontiers in Finance conference (2012), the 3L Finance Research Workshop in Brussels (2013), the Recent Advances in Algo and HF Trading conference at UCL (2013), the 3rd Humboldt-Copenhagen Conference on Financial Econometrics (2013), the Isaac Newton Institute (2013), the City University of Hong Kong Workshop on High-Frequency and Algorithmic Trading (2014), the OMGI Quantitative Investment Seminar at Oxford University (2014), the Cambridge-INET Workshop on Empirical Market Microstructure (2014) and also at Liverpool University and Toulouse Business School. Useful comments were received from all of these audiences. Opinions expressed in this paper are those of the authors and not of Blackrock, Inc. All errors are our own.

†Faculty of Finance, Cass Business School, City University, London. Correspondence: richard.payne@city.ac.uk.

very blunt tool and, at worst, may impose significant costs on those who use low-latency systems to execute efficiently.

Keywords: *Market microstructure; liquidity; transactions costs; asymmetric information; London Stock Exchange; high-frequency trading.*

1 Introduction

The impact of automated trading and quoting by computers on equity price dynamics, execution costs and market stability is firmly under the spotlight. Regulators around the world are looking closely at how algorithmic trading affects the quality of markets and are trying to decide whether and, if so, how market microstructures should be altered so as to limit the scope of computer-generated trading.[1]

Central to this debate is the role of speed in trading.[2] Speed is beneficial for algorithms seeking to execute orders in an agency capacity, in that lower latency implies a greater ability to capture attractive trading opportunities (Biais and Woolley, 2011; Chaboud, Chiquoine, Hjalmarsson, and Vega, 2013). However, it is also argued that speed differentials can facilitate the exploitation of slower traders by faster traders, not due to any innate superior ability but simply by being able to respond to price-relevant information more rapidly (Hasbrouck and Saar (2011) and Biais, Foucault, and Moinas (2013)). Others have suggested that speed may facilitate manipulative trading strategies (Biais and Woolley (2011)). It is not surprising then that regulation based on speed is being discussed. For example, the European Securities and Markets Authority (ESMA) has begun the consultation process for the review of the Markets in Financial Instruments Directive (MiFID II). MiFID II intends to curb HFT and ESMA has proposed a definition of HFT which is based on (i) the close proximity of a firm's server to the trading venue's matching engine; (ii) a connection bandwidth close to the maximum technologically available; and (iii) a high average messaging frequency.[3] If this definition is adopted it suggests that regulation will not be based upon a participant's trading *strategy* but on the participant's *technology*

[1]In the UK, which is the focus of this study, the Department for Business, Innovation and Skills set up a Foresight programme to investigate how computer-generated trading is impacting upon UK markets. See http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading.

[2]It is worth noting that speed differentials are not a feature solely of automated markets. They quite clearly existed in dealerships and floor-trading settings, albeit in less obvious technological form.

[3]An alternative definition proposed by ESMA is based upon identifying firms with a low median lifetime of modified/cancelled orders.

and, in particular, on how quickly the participant can receive and process data and then send a message to an exchange. This is, therefore, regulation based on speed not style.

In the existing academic literature, there is general agreement that certain aspects of low latency trading may be beneficial, for example high frequency market making.[4] Brogaard, Hagströmer, Norden, and Riordan (2013), for example, demonstrate that speed upgrades benefit market makers by improving their ability to control inventory. This in turn reduces transactions costs for all participants and provides investors with enhanced terms of trade. However, due to the previously described concerns regarding manipulation and the role of latency differentials in generating information asymmetries, fast *liquidity-consuming* activities are frequently seen as being harmful and so worthy of further study.

Our paper contributes to this literature by focusing on a specific subset of order book activity defined by speed. More precisely, we study *low-latency liquidity-consuming orders*, which we term fast aggressive trades. In common with most publicly available databases, the trade and quote data we use do not contain any identifiers for the traders involved and thus we cannot relate orders to participants and then to their trading style. Thus we use an alternative mechanism based purely on the reaction speeds of orders.[5] We examine each marketable order in the data and classify it according to the age of the most recent standing limit order against which it executes: 'Slow' trades execute against limit orders that have stood in the book for at least ten seconds and make up around 50% of all trades; 'Immediate' trades execute against limit orders no more than one millisecond old and make up around 10% of all trades; 'fast' trades execute within 100ms and we break them down into three sub-categories in the detailed analysis below. Together, these account for around 20% of trades.[6]

[4]There is still debate over whether these benefits are supplied by low latency traders in less benign market conditions, however.

[5]Our data-based approach to the identification of computer-generated trading resembles the analysis in Hasbrouck and Saar (2011) and Jiang, Lo, and Valente (2013).

[6]Of course, it is not guaranteed that a transaction against a very recently posted limit order is the result of the actions of a low-latency liquidity-consuming algorithm. Just by chance, two orders may be entered within milliseconds of one another and match. This is especially likely in more

The remaining trades that execute between 100ms and 10 seconds are allocated to two 'intermediate' categories. We then compare the characteristics of each speed category with those of slow trades. In particular, we investigate whether trades from faster systems tend to execute at lower cost than those from slower systems, whether they carry more information and how they are timed relative to periods of possibly manipulative activity.

We apply our identification mechanism to data on FTSE 100 stocks electronically traded on the London Stock Exchange, the main UK stock market. The dataset spans Jan-March 2015. Figure 1 highlights some of our key results. It plots average best bid and offer curves prevailing around market buy orders of categories of traders, based on all observations for all stocks in the sample. The quote prevailing immediately before a transaction is timed at -1 on the horizontal axis (which is measured in event time) and the mid-price is normalized to zero at quote -10. The quote at time 0 is that prevailing immediately after the trade has consumed liquidity.

In terms of execution quality, Figure 1 makes clear that relatively fast buys trade at significantly cheaper prices than do slow buys (executing on the dotted line offer curve). Fast trades occur at prices more than 3bp better than equivalent slow trades. This is the key positive aspect of fast trading highlighted in the literature.[7] In the cross-section, execution costs of fast trades are around 50% lower than for slow trades in the most active stocks traded in London, while they are around 40% lower for the least active tercile of large stocks. The benefits of using fast trading techniques in terms of trade execution costs are clearly economically significant.

A second feature of fast trading is that their price impact is, on average, negligible. As Figure 1 shows, once liquidity is consumed by a fast trade, bid, offer and mid-quotes quickly return to almost exactly the same levels that were prevailing ten quotes before the transaction. It appears that fast trades simply consume attractively-

liquid markets where orders, low or high-latency, are frequent. At the other end of the spectrum, a computer algorithm might decide to hit a much older order if other market conditions have changed. Thus our classification scheme is clearly not perfect but in the internet appendix to the paper we demonstrate the robustness of our findings.

[7] Hendershott and Riordan (2012) provide similar evidence for German stocks.

priced liquidity on offer and do not convey any price-relevant information to the market. Slow trades provide a very different picture, and on average result in long-lived movements in both bid and offer prices of almost 4bp.

One might expect there to be some degree of continuity between fast trades of differing speeds. This does appear to be the case for all but immediate trades. Figure 1 show that in the run-up to the trade (times between -8 and -2), the offer curves are (almost) monotonically higher for progressively faster trades. The offer prices immediately prior to the trade tend to bunch up but there is still an ordering that correlates with speed. After the trade (time 0 to 10), the offer curves spread out again and faster trades tend to see faster recovery towards the benchmark time -10 price.

It is striking, however, that *immediate* buys behave very differently from other very fast trades. They are, in fact, much more similar to slow trades. They do not execute at particularly advantageous prices but they do have substantial long-term price impacts. Immediate trades execute against limit orders less than one millisecond old and their associated offer curve is very different from what we will term "ultra-fast" trades, which execute between 1-10ms. However, if we consider "near-immediate" trades, which execute against limit orders more than one millisecond but less than two milliseconds old, we find that these resemble other fast trades, tending to execute against attractively priced offers and having no price impact in the long run. Immediate trades are clearly different from all other fast trades, even those just very marginally slower.

This lack of long-run price impact for most classes of fast trades is in stark contrast with much of the high frequency trading literature which suggests that low latency HFT trades in U.S. markets convey *more* information than slow trades (Brogaard, Hendershott, and Riordan (2012), Kirilenko, Kyle, Samadi, and Tuzun (2011)). This information asymmetry raises concerns about adverse selection costs being imposed on higher latency traders, which are seen as the counterpoint to the execution cost advantages outlined earlier (Biais, Foucault, and Moinas (2013), Brogaard, Hendershott, and Riordan (2012)). In our case, however, it appears that fast trading imposes no adverse selection costs on slow traders. This is not true of immediate trades whose price impact after the trade is substantial. However, such immediate trades are very comparable with our benchmark slow trades in that they bring information to the market and pay substantial transactions costs to execute (despite their speed advantages).

In reconciling our results with those of previous work, it is worth making clear that our speed-based measure is not a measure of HFT. It will capture some of the activities of high frequency traders (HFT) but it also includes many transactions from execution algorithms, among others.[8] We make no claim to be isolating HFT activity, nor do we regard this as a study of pure HFT. Our focus is on characteristics of trading emanating from the fastest systems and thus the possible implications of regulating trading speed, complementing theoretical work by Biais, Foucault, and Moinas (2013), Budish, Cramton, and Shim (2013) and Pagnotta and Philippon (2011) and related empirical work such as Brogaard, Hagströmer, Norden, and Riordan (2013).

Our final contribution is to examine whether counterparties are being induced to transact against fast aggressive traders to the latter's benefit. Much of the regulatory discussion of high frequency trading has focused on possible manipulation of markets by low latency traders. For example, spoofing, layering, and fading (as explained later in the paper) are techniques which may all rely to some extent on speed differentials. We therefore consider in detail both the actions of the counterparties of fast trades and the evolution of markets in the run-up to fast aggressive trades. We find little evidence that our fast trades are manipulative.

The rest of the paper is set out as follows. Our data and the method used to identify fast trades are described in Section 2, together with some descriptive statistics. Section 3 presents our empirical analysis of the performance of fast trades and considers issues relating to manipulation. Section 4 concludes.

[8]The Association for Financial Markets (2010) reports that the estimated market share of all HFT on the London stock exchange in 2010 was 33% and so the 6-25% captured by immediate-fast trading is clearly not the whole story of HFT activity.

2 Data

2.1 Data sources and coverage

We base our analysis on data from the London Stock Exchange provided by Reuters' TRTH data service which we access via the European Capital Markets Cooperative Research Centre. We observe trade and quote data for each stock, with quote data covering the first three prices on each side of the order book. That is, we observe the best three prevailing bid and offer prices and associated depths, together with the price and size of all trades. All data are timed to the millisecond. Increases in depth in the order book represent the entry of new limit orders and strict price and time priority rules allow us to compute the age of any limit order at the time of its execution.[9]

The stocks we consider all trade on the Stock Exchange electronic Trading Service (SETS) electronic limit order book. The SETS order book uses standard trading protocols and is open between 8.30 and 16.30 local UK time. As trading is opened and closed with an auction period, we omit the first and last 5 minutes of the trading day from our analysis. Our starting set of securities are the 100 constituent stocks of the FTSE 100 index. Even though these are the largest listed companies in the UK, activity and market capitalisation vary considerably in the cross-section and so we perform most of our analysis separately on terciles of stocks grouped by average daily trading volume (ADV) computed over each sample.

Based on the trade and best quote data we can compute the age of any limit order at the time of its (partial or complete) execution.[10] This will be the crucial measure used in this paper to identify aggressive computer generated activities, as explained

[9]We only observe the best three prices on each side of the book. New depth outside these prices is not be accurately timed. Such on order is first observed in our data when it becomes (part of) the depth prevailing at the third best price, even if it has been in the book for much longer. The age of such orders at execution will be underestimated.

[10]Executions against multiple standing orders in the limit order book, whether as a result of a market order or a marketable limit order, result in several trade messages. We group these as one event and compute the execution price as the appropriately weighted average execution price.

in the next subsection.

2.2 Identifying the speed of aggressive trades

We define the trade speed of any market or marketable limit order to be the time difference between the entry of that aggressive order and the most recent entry time of a passive limit order that it at least partially executes against.[11] Trades are aggregated so that executions at the same price, with the same direction and at the same millisecond are assumed to be one execution. Note too that the order book may change through entries, cancellations or executions during the relevant interval such that the entry of the limit order and execution need not be consecutive events. This approach has the important advantage that low-latency trades can be identified using public data and without researchers needing to gain access to often highly confidential trader identification data.

Each aggressive trade is then allocated to one of several categories based on speed:

1. 'Immediate' execution, where the trade speed is less than 1 millisecond
2. 'Ultra-fast' execution, where the trade speed is between 1-10ms
3. 10-25ms
4. 25-100ms
5. 100ms-1 second
6. 1 second - 10 seconds
7. 'Slow' trades, where the trade speed is greater than ten second

A few words about our terminology are in order at this point. We will call trades 'slow' if they execute against limit order that were posted more than ten seconds previously

[11] Our definition of fast trades is closely related to that used by Jiang, Lo, and Valente (2013), to identify high frequency trades in the fixed income market.

(i.e. they fall into category 7 in this list). The first four categories represent 'fast' trades and trades in categories 5 and 6 (so executing between 100ms and 10 seconds) are intermediate. In our subsequent analysis, the 'slow' category will usually be our benchmark and we explicitly compare the characteristics of faster trades with those of these slow trades. However, our classification captures variations in fast trades and we also compare characteristics between these. We term the very fastest trades 'immediate' but this is not literally correct. London Stock Exchange data are send to Reuters with a one millisecond time-stamp. We therefore only use millisecond level timings in the analysis. Immediate trades are those which are record in the data as executing at the same time as the limit order (to the nearest millisecond) but in reality they are generated by algorithms with a latency of less than one millisecond.[12] The 'immediate' execution category will turn out to be particularly important. Trades in this category appear different to all other fast trades and in fact look more like slow trades. Because of this, we will spend time comparing 'immediate' trades with those in the next category of 1-10ms which we will term 'ultra-fast'. Finally, we acknowledge that, with the possible exception of the immediate trades, the various categories are obviously partitioned in a rather arbitrary way. However, the conclusions we reach are invariant to alternative categorisations, as explained in more detail below.

We also recognise that our classification scheme is not foolproof. A market order emanating from a high latency (i.e. slow) trader may, by chance, execute against a limit order that was entered only a few milliseconds previously. However, the chances of this happening are slight relative to the probability of such a trade resulting from a low latency trader, particularly for less actively traded stocks.

A related concern is that we might systematically misclassify fast trades under particular market conditions (e.g. when trading activity is particularly heavy). To better understand this issue we compute stock-by-stock correlations between time-series of fast trade proportions (defined as the number of fast trades divided by the total number of trades) and the number of trades, trading volume, and average bid-ask spreads,

[12] Trades in the Reuters TRTH database are timed to the microsecond but this is a spurious level of accuracy.

each computed over five-minute intervals. We then average the correlations across stocks. Correlations are low - the highest is just 0.11 for fast proportions and the number of trades, and the lowest just 0.08 for fast proportions and spreads - suggesting that our FAT arrivals are not particularly strongly related to market conditions. We also examine intraday patterns of the same variables (Figure 2). Trading activity on the London Stock Exchange increases noticeably towards the end of the day, while spreads decline from their initially high levels before plateauing.[13] However, the proportion of fast trades is remarkably constant throughout the day. These same patterns still hold if we split stocks into terciles by average daily trading volume. We are therefore encouraged that our measure of fast trading is not driven by unusual or particular market conditions.

At the other end of the spectrum, we will misclassify trades when a low latency trader decides to hit an order that is minutes old, perhaps due to a public information release or changes in market conditions in a related stock. We cannot measure the probability of such a event occurring but suspect that it might happen frequently in our data. To the extent that this means that our slow trades are contaminated with some faster activity, any differences we find between the characteristics of fast and slow activity are likely to be understatements of the true effects.

Figure 3 plots the average proportion of trades falling into four speed buckets for each stock. Stocks are ordered by daily trading volume (low to high). The first speed bucket in the figure corresponds to the 'immediate' category discussed above, namely execution within one millisecond. While there is cross-sectional variation in the proportion of trades that fall into this category there is no obvious pattern and the proportion oscillates around ten percent. The second bucket combines all trades with execution speeds of 1-100ms. This bucket corresponds to 'fast' (but not immediate) trades, pooling together categories 2-4 above. This bucket constitutes around twenty percent of trades, slightly less for the least active stocks and rising to around one-quarter of all trades for the most active stocks. The third bucket captures

[13]The single peak in trading volume is due to extremely unusual activity on the day that the short sales ban in financial stocks was introduced. Excluding this day has no effect on our results.

intermediate execution speeds of between 100ms and 10 seconds, and also captures around twenty percent of trades for most stocks, rising to over one-third for the most active few stocks. All other trades fall into the fourth bucket which corresponds to the benchmark 'slow' category. Only one-third of trades fall into this category for the most active stocks but this proportion rises steadily to a little over one-half as we consider less active stocks. Nevertheless, for almost all stocks, the category contains more trades than any other. Recall that Table 2 shows that slow trades are also typically larger than faster trades. Slow trades are therefore the most common type of trade and, when they occur, are also larger than the alternatives.

2.3 Descriptive statistics for fast aggressive trades

We first consider the simple cross-section variation in fast trading activity. Table 1 presents summary statistics averaged across stocks grouped into three terciles according to ADV. T1 is the lowest ADV quintile (the least active stocks) and T3 the most actively traded. Many of the basic descriptive statistics are as expected. T1 stocks have the largest mean bid-ask spread at nine basis points and smallest mean trade size of GBP6,400 per transaction. Note that transactions with the exact same time stamp and in the same direction are aggregated into one transaction for these calculations. Though all shares we consider are constituents of the narrowly defined benchmark FTSE100 index and hence are all actively traded in an absolute sense, volume is significantly higher for top tercile stocks.

The final seven columns of Table 1 are perhaps of more interest. These give the percentages of trades that fall into each of our trade speed categories. These do not differ noticeably across terciles and so we focus on T3 figures in our discussion. A little over 9.5% of trades are immediate, in that they at least partially execute against a limit order that is less than one millisecond old. Ultra-fast trades that respond within 10ms make up a further nine percent, and by the time we have accumulated all trades responding within 100ms of the limit order being posted, fast trades account for around 40% of all trades. This fast proportion is slightly higher for more actively traded shares which may, in part, reflect contamination of our speed measure when a coincidental high latency trade is misclassified as a low latency one. It may equally reflect the necessity of participants to trade more quickly in more liquid securities.

Table 2 details the mean trade size broken down by speed category. Trade size is, on average, larger for more active stocks both overall and in each speed category. More importantly, there is a pronounced U-shaped relationship between mean trade size and speed. Mean trade sizes in slow and immediate categories are approximately equal but are around one-third smaller in the 25-100ms speed category. This comparability of results for the very fastest and the very slowest categories will be a recurring theme in the paper.

3 Empirical analysis

3.1 Execution Costs

Speed can be critical for executing trades at the best possible price. Many limit orders are only fleetingly available and high latency traders may not be quick enough to trade against them. It would therefore seem likely that faster traders can execute trades cheaper than slower traders. One contribution of this section is to show just how large the savings from being fast can be. The second contribution is to also demonstrate that speed does not necessarily mean execution cost savings. The very fastest trades which have almost immediate execution have higher execution costs than somewhat slower trades.

We assign an execution cost to each trade. This cost is the distance, in basis points, between the trade price and the mid-quote just prior to the execution of the trade and, as SETS is fully order driven, must equal at least half the bid-ask spread and maybe more if the order is large enough to walk the book.[14] This measure is guaranteed

[14]EDIT: Note that, at the time of our sample, hidden orders were allowed but were very rare in trading these stocks. When encountered, a hidden order could reduce the cost of trading below the observed half spread.

to be positive for buy orders in an electronic market, and for sell orders it is always negative. Thus, we take the negative of the measure for sell trades. Denote that measure for trade t in stock i by $z_{i,t}$. We then regress execution costs on a constant, seven dummy variables to indicate which speed category the trade lies in (denoted $j = 1...7$) and controls for the size of the trade (measured in thousands of GBP), overall stock-level trading volume (measured in millions of GBP) over the fifty trades preceding the current trade, and stock volatility measured as the square root of the sum of squared basis point returns over the preceding fifty trades:

$$z_{i,t} = \sum_{j=1}^{j=7} \beta_{1,i,j} Speed_{i,j,t} + \beta_{2,i} Size_{i,t} + \beta_{3,i} \sigma_{i,t} + \beta_{4,i} Volume_{i,t} + u_{i,t} \tag{1}$$

This model is estimated by a pooled panel regression, both for the full sample of stocks and separately for each ADV-based tercile of our stocks. There is no constant term in the regression since the trade speed dummies saturate the model. All explanatory variables except the trade speed indicator dummies are demeaned and so the coefficients on the speed dummies give the mean execution cost for trades in that category, controlling for trade size, volatility and volume.

Given that SETS is a pure limit order book, these execution cost regressions ask whether, controlling for market conditions, there are systematic differences in bid-ask spreads at times when trades of different speeds execute. Alternatively, do executions in faster categories tend to capture smaller spreads than slower executions? Table 3 presents the results.

Consider first the control variables. As expected, execution costs are higher in more volatile periods and for larger transactions. Conversely, but still as expected, execution costs are lower in high volume periods. The magnitudes of the effects of size and volume decrease as we consider more liquid stocks but the magnitude of the effect of volatility is increasing with average stock liquidity levels.

The column labelled 'Slow' gives the average basis point trading costs for our benchmark category of slow executions (more than ten seconds after the arrival of the preceding limit order against which it executes). These decline across activity terciles in the expected manner. Moving to the key variables in the analysis, the results show that faster trades execute at significantly lower costs than slow trades in all regressions. As we consider progressively faster categories of trade, we observe that execution costs fall. The gain relative to slow trades is around 1.5 basis points for trades in the 25-100ms category and almost 2.5bp for trades in the 1-10ms category when we consider all stocks. Put differently, trades in the 25-100ms category are over 40% cheaper than slow trades, and ultra-fast trades in the 1-5ms category are over 60% cheaper than slow trades. Gains are proportionately larger for more liquid stocks and ultra-fast executions in the most liquid stocks are 72.2% cheaper than equivalent slow trades. Speed can translate into very significant execution cost savings.

Notably, however, transactions costs for the immediate category are *higher* than for other very fast trades. The increase in execution costs for immediate trades relative to ultra-fast 1-10ms trades is economically large 0.9bp for the most liquid stocks and 1.6bp for the least liquid tercile. It is apparent that speed helps to reduce execution costs but the lowest latency trades do not appear to be driven by the need to shave yet more off trading costs. In fact, execution costs of immediate trades are comparable to those of trades in the relatively pedestrian 1-10 second category.

Equation (1) linearly conditions execution costs on trading activity. We proceed to investigate whether the cost savings generated by trade speed change systematically with trading activity by running regressions similar to (1), but for volume-based subsamples of the data. For each stock we identify the quartile of data points with the lowest volume and the quartile with the greatest volume (where volume is defined over the past fifty trades). We then re-estimate equation (1) with stocks double-sorted by ADV and volume (we also pool all stocks and only split by volume). Results are reported in Table 4. The coefficient values show that average costs for the various categories of trade execution speeds only differ slightly across the volume quartiles. On average, as expected, realised spreads are slightly higher in less active periods for all categories. This average effect does not survive when we consider stocks split into terciles so, for example, for the most actively traded stocks coefficients on the speed categories are each lower in low volume periods than in high volume periods. This surprising result is, however, driven by significant changes in the coefficients on the conditioning variables. More importantly, the results split by ADV and volume reveal the same patterns as found previously - slow trades are the most expensive to execute and transactions costs fall for each progressively faster category of trades. The exception is the immediate category. Immediate traders still resemble slow traders in that they pay higher execution costs irrespective of activity levels and despite their apparent speed advantage.

These results suggest that fast algorithms can offer very significant advantages in execution cost over slow trades, **especially for less actively traded stocks or for less active trading periods**. The speed advantage offered by fast trading translates directly into access to much better prices. When trading and quoting is less continuous (due to low natural trading interest or high tick sizes), the fastest algorithms can take good prices as they appear and before other participants (human or high latency algorithms) can react. Thus, as one would expect, speed directly translates to lower costs of trading. However, the very fastest category of trades, those that execute within a millisecond of a new limit order hitting the book, do not trade to gain more savings. Execution costs of these trades are typically higher on average than all other categories of fast trades. We conclude that traders willing to pay substantial amounts of money to attain immediate execution gain in ways other than reduced execution costs.

3.2 Information content

Computer-based trading, and in particular, high frequency trading is often characterized as predatory and of no social value. However, much of the empirical literature suggests that high frequency trading brings, on average, more information to the market than does slower trading activity (Brogaard, Hendershott, and Riordan (2012), Brogaard, Hagströmer, Norden, and Riordan (2013)). As Martinez and Rosu (2013) suggest, this may be due to HFTs receiving or being able to react to information about payoff-relevant events ahead of the rest of the market. These findings help to motivate theoretical models with information asymmetries (Martinez and Rosu (2013), Foucault, Hombert, and Rosu (2013)) and lead to arguments in which computer based trading is beneficial as it leads to more efficient markets. The downside is that if fast traders are relatively more informed then trading against them will be costly. If counter-parties are adversely selected by fast traders then they may widen spreads and reduce depths, increasing transactions costs for all aggressive traders.

To investigate the extent to which fast trades contribute more information to markets than slow trades, we measure the returns around trades and investigate whether they are systematically different for trades of differing speeds. We estimate the equation below by pooled OLS;

$$r_{i,t-k,t+k} = \sum_{j=1}^{j=7} \beta_{1,i,j} Speed_{i,j,t} + \beta_2 Size_{i,t} + \beta_3 \sigma_{i,t} + \beta_4 Volume_{i,t} + \epsilon_{i,t} \quad (2)$$

where $r_{t-k,t+k}$ is initially defined to be the midquote return between k order events before the transaction of interest and k order events after the trade (with sign swapped if the current trade is a sell). $Speed_{i,j,t}$ is a dummy indicating which speed category the relevant trade falls into (which again saturate the model so there is no constant term included), and the other variables are controls for trade size ($Size_t$), market volatility (σ_t) and market volume ($Volume_t$) as defined in Section 3.1.

Panel A of Table 5 gives the results of the preceding regression based on k=10 for the full sample and the three ADV-based activity categories. That is, we measure the price impact of trades from the mid-quote prevailing ten order book events prior to the transaction of interest until ten events later ('post trade'). Results for the conditioning variables suggest that larger transaction sizes increase price impacts, while greater levels of trading activity decrease price impacts, as expected. Recent price volatility is not significant.

More importantly, we note that the price impact of a slow trade is over 3bp for the most active stocks, rising to 4bp for the least liquid tercile of stocks. When we consider faster trades three clear results emerge. First, trades in all other speed categories have lower price impacts, on average, than trades in our benchmark slow category. Average price impacts drop rapidly once we consider faster trades and even relatively slow trades in the 1-10 second category have price impacts of less than one basis point. Second, relatively fast trades in the 1ms-1 second speed categories have either an insignificant or even negative price impact over this interval. The only statistically significantly positive price impact estimates are for trades in the ultra-fast 1-10ms category but even the largest of these are economically tiny at one-fifth of a basis point. Most fast trades then appear to contain effectively no information. The third finding is an exception to this statement. Immediate trades do have large, positive and statistically significant price impacts. The magnitudes of their average price impacts are not quite as large as those of slow trades but it is clear that this category of fast trades behaves very differently from all other fast trades.

Panel B of the table considers price impact measured using the price on the side of the market *not* directly affected by the deal. That is, using offer-side returns if the trade is a market sell and bid-side returns for an aggressive buy. The analysis of opposite side returns is designed to remove the mechanical liquidity-consumption effect of a trade felt by using mid-quote (or same-side) returns. The price impacts of all categories of trades are not much different from midpoint returns for all three ADV terciles, suggesting that price impacts reported in Panel A do not contain important liquidity effects. This is perhaps not surprising given our focus on FTSE100 index constituent stocks.[15]

As in the previous section, we may be concerned that we have not conditioned on market activity levels adequately by simply including recent trading volume in equation (2). We therefore again double sort the data by ADV and stock-level volume

[15]These results tally in their main features with the simple graphical analysis of quotes around fast and slow trades presented in Figure 1. The magnitudes are not the same using the two methods, but it should be noted that the regression analysis conditions out the effects of several variables that are present in Figure 1.

terciles and re-estimate equation (2). The results reported in Table 6 focus on opposite side price impacts and should be compared with those given in Panel B of Table 5. The price impact of all trades is higher in the low volume subsample. Indeed, all speed categories appear to have positive price impact during less active periods, and almost all are both statistically and economically significant. In the high volume subsample, price impacts are lower and are significantly negative for fast trades. Nevertheless, despite these differences in price impacts across volume quartiles, the key finding that immediate and slow trades have much larger price impacts than fast (and intermediate) trades remains, no matter how the samples are created.

Our findings both contrast and complement the HFT literature. In part, this is likely because of our focus on speed measures whereas measures of HFT usually focus on trading strategy and trader type. Speed is obviously a tool used by HFT, but it is also used by many other agents. Baron, Brogaard, and Kirilenko (2014), for example, note specifically that liquidity consuming HFT earn profits at the expense of other market participants. It is also worth noting that their results are from US data, and the complicated market architecture of the US as compared to the relatively simple UK setting might generate profit opportunities for aggressive HFTs that are available only for the very fastest in our UK data. Regardless of the source of the differences, however, our results suggest that most fast trading is systematically less well informed than slower trading.

3.3 Is there a dark side to fast aggressive trading?

So far our analysis has suggested that fast aggressive traders bargain hunt. Algorithms lie in wait for attractively-priced limit orders, picking them off within fractions of a second. The counterparties to these transactions achieve better prices than they would through using market orders and with quite high probabilities of being filled. As portrayed, this is the natural meeting of traders in the marketplace where both sides are happy to transact at the executed price. However, it is known that certain market conditions can lead traders to place attractively-priced limit orders (Biais, Hillion, and Spatt, 1995; Ranaldo, 2004) and commentators and regulators are concerned that algorithms may be manipulating conditions to induce such orders which can then be picked off through a FAT. Algorithmic traders are regularly associated with dubious practices such as layering or spoofing, designed to mislead other participants as to supply and demand for an asset, such that they make trading decisions that can be exploited by fast algorithms. This is the perceived dark side of computer based trading. In this section, we examine activity in the market around FAT executions, looking for evidence that the FAT was exploiting manipulative practices.

We focus on a particular type of manipulation that is often discussed called 'spoofing' (also known as 'layering'). Suppose a fast trader wishes to sell. By placing several limit orders at or close behind the best bid, the fast trader gives the impression of great depth on the buy side. The hope is that other traders will now submit buy limit orders ahead of the prevailing bid, either in order to gain priority or because they feel the imbalance in the order book is indicative of a likely price rise. The fast trader then uses her speed advantage (i) to execute against the new attractively priced bid, and (ii) to remove all of the spurious bid orders before they are executed.[16]

3.3.1 Do spoofing events increase the probability of fast trading?

Our empirical strategy defines spoofing episodes using our order level data. We define a limit buy side spoofing event to have occurred for a stock if net quote activity, defined as net new limit buy orders minus net new limit sell orders, in a 1 second interval is greater than five times its stock-specific standard deviation.[17] Limit sell side spoofing events are defined analogously. The counts of spoofing events for ADV-ordered stocks over the sample period are shown in Figure 4, where bid-side spoofing columns are above the horizontal, and offer-side spoofing below the horizontal. On average there are around 20,000 spoofing events per stock, with a tendency for there

[16]The fast trader will also be able to exploit her speed while waiting for an attractively priced bid to arrive as, if the market begins to move at this time, she can quickly cancel her spoofing orders before they are executed.

[17]Net new limit buy orders in a 1 second interval equal the number of new buy orders submitted less the number of buy orders cancelled. Net new sell orders is calculated similarly.

to be more events for less active stocks.

We then regress the number of trades within an interval of a specific type (fast or slow, buys or sells) on lagged trades of all type (to control for correlations between order flows); lagged returns, the lagged bid-offer spread, the lagged total number of orders entered ($q_{i,t-1}$), and the lagged net number of orders entered ($netq_{i,t-1}$) to control for market conditions; and indicator variables taking the value unity if buy or sell side spoofing has been identified ($BuySpoof_{i,t-1}$ or $SellSpoof_{i,t-1}$). These final two terms are our variables of interest and their coefficients reveal the impact of spoofing events on the number of trades subsequently observed. Thus for fast aggressive buy trades ($BuyFast$) we run the following regression:

$$\begin{aligned} BuyFast_{it} =& \alpha + \beta_1 BuyFast_{i,t-1} + \beta_2 SellFast_{i,t-1} + \beta_3 BuySlow_{i,t-1} + \beta_4 SellSlow_{i,t-1} \\ & + \beta_5 netq_{i,t-1} + \beta_6 ret_{i,t-1} + \beta_7 q_{i,t-1} + \beta_8 spread_{i,t-1} \\ & + \beta_9 BuySpoof_{i,t-1} + \beta_{10} SellSpoof_{i,t-1} + \epsilon_t \end{aligned} \tag{3}$$

If fast aggressive buy orders are exploiting a manipulative trading strategy we would expect them to follow periods of sell side spoofing, such that β_{10} would be significantly above zero.

All variable are measured at a one second frequency and all explanatory variables except the spoofing indicators are demeaned before the regression is run. Table 7 gives the results, showing coefficient estimates from fixed effect panel estimations (with heteroskedasticity robust standard errors). Note that coefficients on all right-hand side variables except those on the four lagged trade variables are multiplied by 1000 to improve legibility.

The control variables are all statistically significant and generally have the expected effects. β_1 to β_4 are all positive, suggesting that we observe more trades in both directions when markets are active. Wider spreads lead to less fast trading on both sides of the market (β_8). More (gross) quoting activity leads to more fast trades (β_7), while a higher number of new buy limit orders relative to new limit sell orders leads to less fast buying but more fast selling (β_5). These same patterns of effects are also found if we use slow trades as dependent variables (results not reported but available on request).

If fast trades are manipulative, we would expect the number of fast buys (sells) in interval t to be positively related to sell-side (buy-side) spoofing events in interval $t-1$. There is some evidence of this occurring for the largest ADV quintile of stocks, in that for those stocks the coefficient on buy (sell) spoofing in the fast sell (buy) regression is twice as large as that on the sell (buy) spoofing regressor. These relationships are highly statistically significant and relatively large in magnitude - the coefficients on the relevant spoofing indicator variables are much larger than the constant terms in these regressions. Thus, for the most active stocks, the number of FAT observed following a spoofing episode is more than twice the number that would be expected when all other explanatory variables are set to their average values.

However, we do not wish to over-emphasise these results. The goodness of fit statistics in all of these regressions are very low and while the effect of a spoofing episode on the number of fast trades observed is large relative to an unconditional benchmark, the economic effects on trade counts are still extremely small. Nonetheless, for the most active stocks, these findings suggest that spoofing-based manipulation may result in slightly more FAT.[18]

For all other stocks, and for both the fast buys and fast sells, the effects of buy and sell spoofing are almost completely identical and have the same sign. Thus both fast buys and sells are somewhat more likely when there is extreme quoting behaviour on one side of the market. This does not appear to be evidence of manipulation. It may be more simply regarded as an indication that when markets are operating extremely quickly (on either side), the probability of fast trades increases.

[18]We emphasise too that we have no evidence that fast traders are responsible for the spoofing episodes. They may merely be using their speed advantages to react ahead of the competition to such events.

3.3.2 What happens to depth around fast and slow trades?

In the previous subsection we defined spoofing measures using rates of (net) new order entry on the buy and sell sides of the book. These measures were then used to predict fast trading activity. In this subsection we ask a slightly different question. Taking the occurrence of a trade as an event, we measure the changes in order book depth before and after so as to get a sense of what changes in book conditions are driving trades.

This analysis is done separately for fast and slow trades and separately for depth at the best bid and offer and at the first five levels of the order book. We measure depth changes over 5 events pre and post trade and depth changes are expressed as a proportion of average daily volume.[19] Finally, we distinguish depth changes on the same side of the order book as the trade as hit (e.g. the limit sell side for an aggressive buy trade) from those on the opposite side from which the trade has hit (e.g. the limit buy side for an aggressive buy).

Tables 8 and 9 give the results of these measured depth changes (with cross-sectional standard errors in brackets). Let us focus first on Table 8, which presents changes in same side depth. Pre-trade, for both fast and slow trades, depth changes are economically very small (less than one hundredth of 1% of ADV). For slow trades, the same is true post-trade. However, for FAT, post-trade we see consistent and sizeable increases in depth, especially for the least active stocks in quintiles 1 and 2. This would seem to be very rapid replacement of the depth consumed by the FAT itself.

Table 9 gives similar results for the opposite side of the market (i.e. the side of the order book that the trade did not directly affect). For fast trades, there is clear evidence of depth increases pre-trade and similarly sized depth reductions post-trade. These effects appear across all stock quintiles and in both best depth and depth across the first 5 levels. The depth changes around slow trades are much smaller than for

[19] When defining the 5 events before a trade, we exclude the event immediately prior to the current trade as this is likely to be the limit order entry that triggers a fast trade.

FAT, although they tend to be in the same direction.

Putting these results together, do they paint a picture of FAT as manipulative activity? If anything, the results indicate quite the opposite.

If we believe the spoofing story, we would expect a fast buy, for example, to be preceded by an increase in the limit sell side depth, initiated by the spoofer and designed to induce the entry of a new generously-priced limit sell order which the fast buyer then lifts. Subsequent to the trade, the aggressive buyer engaged in the spoofing strategy would cancel the limit sell orders that she had added, leading to a reduction in limit sell side depth post-trade. Our results, however, provide no evidence that fast trades tend to be preceded by run-ups in same side depth which is then removed post-trade. In fact, post-trade we see same side depth rise.[20] Evidence for spoofing is thus weak.

The pre- and post-trade changes in depth for the opposite side of the limit order book are stronger and more significant. The result that both fast and slow trades tend to be preceded by increases in depth on the opposite side of the market might be interpreted as a crowding-out effect as in Foucault (1999). That is, when the buy side of the order book becomes very deep, for example, a trader is more willing to submit a market buy, pay the spread and jump the queue. There is an asymmetry, however, in that fast traders seem to react more strongly to crowding out, especially for less liquid stocks. The result that opposite side depth decreases quite strongly after fast trades is not obviously consistent with the crowding out story, as it is not clear why a trader jumping a limit buy queue would cause others to relinquish their position in that queue of limit buys. An alternative interpretation of these results is that it is the fast traders who are being spoofed. They are aggressively buying when the limit buy side of the order book becomes heavy but after their trade the weight on the limit buy side of the book that induced them to trade disappears. Perhaps this is due to strategic liquidity suppliers exploiting badly designed fast algorithms.

[20] We have investigated the depth changes over different event windows pre-trade with no qualitative change in the results.

3.4 Why are immediate trades so different to the merely ultra-fast?

The analysis of both transactions costs and price impacts suggest that immediate trades appear quite different from other fast trades and, if anything, look more like slow trades. In this section we refine our analysis and compare the properties of immediate trades with those of trades that are nearly but not quite immediate. Specifically, we now split trades in our previous ultra-fast (1-10ms) category into two groups: 'near-immediate' trades that execute 1-2ms after the posting of the limit order, and other ultra-fast trades that execute in the 2-10ms interval.

Estimating a version of equation 1 with the addition of a near-immediate category reveals that there remains a large difference between immediate trades and even near-immediate ones. Further, near-immediate trades are indistinguishable from other ultra-fast trades. For example, when we pool all stocks into one panel regression, the coefficient on immediate trades is 2.996, the coefficient on near-immediate trades is 1.586 and that on other ultra-fast trades is 1.622. Even when we split stocks into ADV terciles this same pattern is revealed.

Results from estimating the price impact equation 2 are only marginally more nuanced. Pooling all stocks into one panel and using midquote returns, the price impact of an immediate trade is 2.872, the price impact of a near-immediate trade is 0.183 and that of other ultra fast trades is 0.029. While near-immediate trades have a much larger positive price impact than other ultra fast trades, economically speaking the magnitude is still tiny at less than one-fifth of a basis point.[21] Again, the same results hold when we examine stocks split by terciles.

Refining the analysis reveals that immediate trades are different from all other type of fast trades, even ones that are only very slightly slower. Transactions costs monotonically fall with execution speed until execution becomes immediate. At this point,

[21] The point estimate for near-immediate trades is also only marginally significant with a t-statistic of just 1.665 though this is due to the relatively few observations falling into this category.

execution costs jump dramatically, rising to the levels seen for trades whose speed is measured in seconds. The information content of fast trades is at best economically tiny and may be statistically significantly negative, though still of small magnitude. This is also true even for near-immediate trades. But the information content of an immediate trade is positive, economically large (averaging around 3bp) and of similar magnitude to slow trades.

4 Conclusions and policy implications

Much of the debate regarding high frequency and algorithmic trading centres on speed. We divide trades in London Stock Exchange equities into two subsets defined by speed. We show that fast trades execute more cheaply than do slow trades, in the sense that effective spreads are significantly smaller. However, fast trades contain almost no information, unlike slow trades. Slow trades are associated with positive price impact, fast trades with zero price impact. Finally, we find very little evidence that fast traders manipulate markets in their favour. If anything, there is more clear evidence that in less liquid stocks some fast traders are themselves be being manipulated.

At least some of the regulatory discussion regarding computer-based trading has suggested the imposition of limits to speed (e.g. minimum resting times) or regulation entities based on the speed of their technology. Our analysis suggests that fast traders are not necessarily the villains that this intuition relies upon as they are, on average, rather innocuous. In our data, speed brings cheap execution but speedier traders are neither fundamentally informed nor malicious. One way to interpret these results is to regard fast trading as being generated, in the main, by activity from sell-side execution algorithms. As such, regulating based on speed (i.e. technology) is, we feel, a poor substitute for regulating trading strategies. In our view, the focus should be squarely on the latter.

References

Baron, M., J. Brogaard, and A. A. Kirilenko, 2014, "Risk and Return in High Frequency Trading," *Available at SSRN 2433118*.

Biais, B., T. Foucault, and S. Moinas, 2013, "Equilibrium Fast Trading," TSE Working Papers 13-387, Toulouse School of Economics (TSE).

Biais, B., P. Hillion, and C. Spatt, 1995, " An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50(5), 1655–89.

Biais, B., and P. Woolley, 2011, "High frequency trading," *Manuscript, Toulouse University, IDEI*.

Brogaard, J., B. Hagströmer, L. L. Norden, and R. Riordan, 2013, "Trading fast and slow: Colocation and market quality," *Manuscript*, 7, 24.

Brogaard, J., T. J. Hendershott, and R. Riordan, 2012, "High-Frequency Trading and Price Discovery," *Review of Financial Studies*, forthcoming.

Budish, E., P. Cramton, and J. Shim, 2013, "The high-frequency trading arms race: Frequent batch auctions as a market design response," *Manuscript*, 6.

Chaboud, A., B. Chiquoine, E. Hjalmarsson, and C. Vega, 2013, "Rise of the machines: Algorithmic trading in the foreign exchange market," *Journal of Finance, Forthcoming*.

Foucault, T., 1999, "Order flow Composition and Trading Costs in a Dynamic Limit Order Market," *Journal of Financial Markets*, 2, 99–134.

Foucault, T., J. Hombert, and I. Rosu, 2013, "News Trading and Speed," Les Cahiers de Recherche 975, HEC Paris.

Hasbrouck, J., and G. Saar, 2011, "Low-Latency Trading," *SSRN eLibrary*.

Hendershott, T., C. M. Jones, and A. J. Menkveld, 2011, "Does Algorithmic Trading Improve Liquidity?," *Journal of Finance*, 66, 1–33.

Hendershott, T. J., and R. Riordan, 2012, "Algorithmic Trading and the Market for Liquidity," *Journal of Financial and Quantitative Analysis*.

Jiang, G. J., I. Lo, and G. Valente, 2013, "High Frequency Trading in the US Treasury Market: Evidence around macroeconomic news announcements," .

Kirilenko, A. A., A. P. S. Kyle, M. Samadi, and T. Tuzun, 2011, "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market," *SSRN eLibrary*.

Martinez, V., and I. Rosu, 2013, "High Frequency Traders, News and Volatility," *SSRN eLibrary*.

Pagnotta, E., and T. Philippon, 2011, "Competing on speed," working paper, National Bureau of Economic Research.

Ranaldo, A., 2004, "Order aggressiveness in limit order book markets," *Journal of Financial Markets*, 7(1), 53 – 74.

Table 1: Summary statistics for market data: stocks grouped into ADV terciles

	Spread	Trade Size	Volume	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow
All	7.592	9.642	24.946	0.091	0.090	0.078	0.054	0.080	0.152	0.455
T1	8.833	6.399	8.059	0.084	0.085	0.073	0.046	0.066	0.131	0.514
T2	8.393	10.149	16.277	0.094	0.093	0.080	0.053	0.073	0.137	0.470
T3	5.551	12.377	50.502	0.096	0.093	0.081	0.061	0.101	0.187	0.380

Notes: each row gives equally weighted mean values of data for the full sample of all stocks or for a particular tercile of stocks, grouped by ADV (i.e. for each ADV group we take the stock-level averages of our variable and then create an equally weighted cross-sectional average. Note that all stock level averages are trimmed means, where the top and bottom percentiles of the distribution have been discarded for robustness.) T1 is the lowest ADV portfolio and T3 the highest ADV portfolio. The first three columns of data give mean bid-ask spreads in basis points, mean trade size in thousands of GBP, and mean volume traded per day in millions of GBP. The columns headed $FAT_{<1}$ to FAT_{1s-10s} give the percentages of Fast Aggressive Trades in each of the various buckets as described in the text. The final column gives the percentage of slow aggressive trades.

Table 2: Mean trade size by speed category

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow
All	11.113	8.421	7.696	7.106	7.361	8.620	10.979
T1	7.438	5.701	5.257	4.792	4.889	5.545	7.076
T2	11.699	8.826	8.060	7.403	7.713	8.960	11.776
T3	14.203	10.736	9.769	9.124	9.480	11.355	14.085

Notes: Each row gives equally weighted mean trade size measured in thousands of GBP by speed category for the full sample of all stocks or for a particular tercile of stocks, grouped by ADV (i.e. for each ADV group we take the stock-level averages of our variable and then create an equally weighted cross-sectional average. Note that all stock level averages are trimmed means, where the top and bottom percentiles of the distribution have been discarded for robustness.) T1 is the lowest ADV portfolio and T3 the highest ADV portfolio. Trades with the same time stamp and direction are aggregated into one transaction.

Table 3: Realised spread: Panel regressions

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	2.995	1.619	1.784	2.407	2.669	3.161	4.183	0.388	0.004	-0.004	0.595
	[17.376]	[11.294]	[13.770]	[20.216]	[23.087]	[26.890]	[33.134]	[31.745]	[10.225]	[-14.959]	
T1	3.932	2.336	2.536	3.277	3.553	4.018	5.187	0.587	0.002	-0.008	0.608
	[18.359]	[13.843]	[16.855]	[23.235]	[25.903]	[29.069]	[35.006]	[26.601]	[7.907]	[-17.064]	
T2	3.437	1.805	1.976	2.660	2.939	3.558	4.777	0.386	0.002	-0.003	0.610
	[17.595]	[11.999]	[14.340]	[20.128]	[22.234]	[26.023]	[32.310]	[30.537]	[8.440]	[-12.815]	
T3	1.617	0.717	0.840	1.285	1.516	1.907	2.586	0.192	0.007	-0.001	0.567
	[16.174]	[8.039]	[10.115]	[17.285]	[21.125]	[25.577]	[32.086]	[38.098]	[14.327]	[-14.997]	

Notes: Stocks are grouped into ADV based terciles. A pooled panel regression is run of realised spreads, measured in basis points, on a set of fast trade indicators, a slow trade indicator, trade size (measured in thousands GBP), stock volatility (measured as the square root of the sum of squared basis point returns over the last 50 trades) and trading volume (measured as millions of GBP traded over the last 50 trades). The final column gives the regression R^2. Each regression has at least 1 million observations.

Table 4: Realised spread: fixed effect panel regressions

Panel A: high volume quartile

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	2.797	1.234	1.358	1.871	2.127	2.751	4.025	0.314	0.005	-0.002	0.590
	[8.330]	[4.179]	[4.972]	[7.520]	[8.879]	[10.944]	[15.032]	[20.208]	[9.544]	[-5.031]	
T1	3.051	1.258	1.431	2.036	2.294	2.830	4.170	0.473	0.004	-0.004	0.609
	[6.521]	[2.984]	[3.819]	[6.051]	[7.122]	[8.665]	[12.371]	[17.252]	[8.922]	[-4.299]	
T2	3.218	1.431	1.545	2.069	2.397	3.183	4.716	0.310	0.005	-0.002	0.615
	[7.548]	[3.605]	[4.292]	[6.141]	[7.364]	[9.441]	[13.331]	[19.133]	[9.869]	[-4.305]	
T3	2.121	1.011	1.096	1.508	1.689	2.242	3.188	0.160	0.007	-0.001	0.545
	[10.923]	[5.947]	[6.805]	[10.369]	[12.152]	[14.725]	[19.395]	[24.237]	[9.842]	[-6.488]	

Panel B: low volume quartile

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	3.456	2.296	2.499	3.343	3.558	3.973	4.758	0.527	0.004	-0.018	0.602
	[9.444]	[7.397]	[8.546]	[12.046]	[13.578]	[15.266]	[17.607]	[16.891]	[6.964]	[-8.992]	
T1	4.903	3.536	3.751	4.816	5.040	5.458	6.469	0.817	0.001	-0.035	0.618
	[11.403]	[10.021]	[11.146]	[14.399]	[16.162]	[17.892]	[20.841]	[14.517]	[4.734]	[-10.616]	
T2	3.911	2.491	2.734	3.664	3.859	4.381	5.305	0.527	0.002	-0.016	0.613
	[10.160]	[8.060]	[9.339]	[12.595]	[13.647]	[15.353]	[17.840]	[16.182]	[5.520]	[-8.371]	
T3	1.555	0.860	1.012	1.549	1.776	2.082	2.501	0.238	0.008	-0.005	0.577
	[6.770]	[4.111]	[5.154]	[9.145]	[10.926]	[12.552]	[14.139]	[19.976]	[10.638]	[-7.989]	

Notes: Stocks are grouped into ADV based terciles. A pooled panel regression is run of realised spreads, measured in basis points, on a set of fast trade indicators, a slow trade indicator, trade size (measured in thousands GBP), stock volatility (measured as the square root of the sum of squared basis point returns over the last 50 trades) and trading volume (measured as millions of GBP traded over the last 50 trades). The final column gives the regression R^2. Each regression has at least 1 million observations.

Table 5: Tick time price impact regressions: -10 to 10 trades

Panel A: midquote returns

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	2.872	0.040	-0.270	-0.551	-0.338	0.701	3.805	0.257	0.021	-0.021	0.079
	[25.476]	[3.139]	[-0.351]	[-2.896]	[-0.072]	[16.401]	[70.965]	[6.786]	[1.070]	[-7.806]	
T1	3.721	0.206	-0.137	-0.645	-0.361	0.781	4.073	0.281	0.020	-0.024	0.069
	[18.477]	[1.922]	[-0.124]	[-2.361]	[-0.907]	[7.927]	[41.869]	[3.862]	[1.459]	[-5.017]	
T2	2.971	-0.166	-0.513	-0.692	-0.518	0.530	3.920	0.329	0.013	-0.026	0.063
	[19.806]	[-0.229]	[-2.261]	[-3.371]	[-2.174]	[8.055]	[49.272]	[6.899]	[1.216]	[-5.889]	
T3	1.923	0.081	-0.159	-0.315	-0.134	0.793	3.420	0.159	0.029	-0.014	0.104
	[38.144]	[7.265]	[-1.331]	[-2.955]	[-2.865]	[33.220]	[121.754]	[9.596]	[0.537]	[-12.512]	

Panel B: opposite side quote returns

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	2.953	0.171	-0.163	-0.470	-0.263	0.834	3.945	0.231	-0.110	-0.027	0.065
	[23.142]	[3.812]	[0.547]	[-1.942]	[-0.687]	[15.532]	[62.674]	[5.174]	[-1.406]	[-8.740]	
T1	3.896	0.436	0.056	-0.474	-0.193	1.010	4.300	0.234	-0.057	-0.035	0.055
	[17.088]	[2.989]	[0.973]	[-1.132]	[-0.204]	[8.365]	[36.803]	[2.783]	[-1.032]	[-6.454]	
T2	3.012	-0.053	-0.414	-0.638	-0.487	0.652	4.063	0.309	-0.046	-0.031	0.051
	[17.885]	[-1.065]	[-1.235]	[-2.563]	[-1.576]	[8.133]	[43.515]	[5.310]	[-1.071]	[-6.375]	
T3	1.950	0.130	-0.131	-0.297	-0.110	0.841	3.472	0.149	-0.227	-0.016	0.089
	[34.453]	[7.381]	[-1.904]	[-2.130]	[-3.432]	[30.097]	[107.704]	[7.428]	[-2.113]	[-13.392]	

Notes: Stocks are grouped into ADV based terciles. For each stock in each group, we run a linear regression of returns around trades on a set of fast trade indicators, a slow trade indicator, trade size (measured in thousands of GBP), stock volatility (measured as the square root of the sum of squared basis point returns over the last 50 observations/1000) and trading volume (measured in thousands of GBP traded over the last 50 observations). Returns are measured from the observation just preceding the current trade to 10 observations afterwards. We use either the midquote to compute returns or the level of the quote on the opposite side of the market (i.e. the bid quote if the trade at the current observation is a market buy). An observation is recorded every time there is a trade or a quote update. The sign of returns around sell trades is changed so that all trades should have positive price impacts. The table presents mean coefficients and t-statistics across the stocks in each group was well as the mean R^2.

Table 6: Tick time opposite side price impact regressions: -10 to 10 trades: high and low volume quartiles

Panel A: high volume quartile

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	2.604	-0.450	-0.805	-0.996	-0.679	0.278	3.343	0.149	0.708	-0.007	0.062
	[8.825]	[-0.560]	[-1.892]	[-2.346]	[-1.205]	[3.338]	[18.882]	[2.165]	[0.632]	[-1.318]	
T1	3.755	0.010	-0.386	-0.873	-0.293	0.704	3.908	0.149	0.237	-0.013	0.054
	[7.431]	[0.000]	[-0.776]	[-1.391]	[-0.382]	[2.329]	[11.934]	[0.983]	[0.408]	[-1.457]	
T2	2.398	-1.063	-1.437	-1.484	-1.268	-0.215	3.142	0.206	0.682	-0.006	0.055
	[6.572]	[-1.589]	[-2.598]	[-2.669]	[-1.984]	[0.904]	[12.192]	[2.543]	[0.462]	[-0.441]	
T3	1.658	-0.296	-0.592	-0.630	-0.476	0.346	2.981	0.094	1.206	-0.003	0.079
	[12.471]	[-0.091]	[-2.301]	[-2.977]	[-1.249]	[6.782]	[32.518]	[2.971]	[1.026]	[-2.057]	

Panel B: low volume quartile

	$FAT_{<1}$	FAT_{1-10}	FAT_{10-25}	FAT_{25-100}	FAT_{100-1s}	FAT_{1s-10s}	Slow	Size	Volatility	Volume	R^2
All	4.162	1.569	1.281	0.765	0.979	2.208	5.479	0.447	-0.035	-0.210	0.080
	[10.379]	[4.388]	[3.343]	[1.655]	[3.112]	[8.576]	[21.262]	[3.768]	[-0.771]	[-5.172]	
T1	5.569	2.170	1.877	1.098	1.334	2.756	6.311	0.532	-0.068	-0.289	0.065
	[7.822]	[3.342]	[2.836]	[1.487]	[2.296]	[5.268]	[12.981]	[2.578]	[-0.515]	[-4.203]	
T2	4.525	1.758	1.434	0.942	1.099	2.310	5.931	0.552	-0.048	-0.262	0.062
	[8.477]	[3.715]	[3.064]	[1.780]	[2.634]	[6.139]	[15.523]	[3.542]	[-0.917]	[-5.004]	
T3	2.392	0.779	0.532	0.255	0.504	1.559	4.197	0.258	0.010	-0.079	0.113
	[14.837]	[6.106]	[4.131]	[1.698]	[4.405]	[14.321]	[35.280]	[5.184]	[-0.880]	[-6.309]	

Notes: Stocks are grouped into ADV based terciles. For each stock in each group, we run a linear regression of returns around trades on a set of fast trade indicators, a slow trade indicator, trade size (measured in thousands of GBP), stock volatility (measured as the square root of the sum of squared basis point returns over the last 50 observations/1000) and trading volume (measured in thousands of GBP traded over the last 50 observations). Returns are measured from the observation just preceding the current trade to 10 observations afterwards. We use either the level of the quote on the opposite side of the market to calculate returns (i.e. the bid quote if the trade at the current observation is a market buy). An observation is recorded every time there is a trade or a quote update. The sign of returns around sell trades is changed so that all trades should have positive price impacts. The table presents mean coefficients and t-statistics across the stocks in each group was well as the mean R^2.

Table 7: Spoofing panel regressions, fast buys and fast sells orders

Panel A: fast buys

	Const	FastBuys(-1)	FastSells(-1)	SlowBuys(-1)	SlowSells(-1)	netQuotes	Returns	Quotes	Spreads	BuyLayer	SellLayer	R^2
Q1	-0.0372	0.0256	0.0048	0.0048	0.0043	-0.2610	-0.2532	0.3236	-0.0005	0.8701	0.7874	0.0015
	[-27.7595]	[13.1497]	[7.0739]	[16.2727]	[27.6636]	[-8.8498]	[-0.8501]	[19.4862]	[-64.5179]	[15.8869]	[11.7198]	
Q2	-0.1251	0.0373	0.0090	0.0070	0.0075	-0.3297	0.0282	0.5863	-0.0009	1.8300	1.7089	0.0035
	[-55.5096]	[15.2168]	[15.3174]	[35.2361]	[44.0420]	[-10.0839]	[2.5704]	[28.0498]	[-31.3359]	[20.8382]	[15.3403]	
Q3	-0.2523	0.0453	0.0140	0.0085	0.0079	-0.3023	0.1158	0.6788	-0.0028	2.6900	2.5806	0.0049
	[-71.7465]	[26.1878]	[12.2563]	[41.1761]	[42.5119]	[-10.3503]	[2.8371]	[30.4267]	[-37.3085]	[20.5345]	[17.4687]	
Q4	-0.5427	0.0427	0.0141	0.0103	0.0100	-0.3139	0.0107	0.5866	-0.0003	3.7169	3.7362	0.0062
	[-113.7972]	[30.3625]	[25.5805]	[77.3919]	[75.1853]	[-15.3666]	[8.0512]	[31.1206]	[-26.3971]	[25.4523]	[22.0233]	
Q5	-1.7449	0.0598	0.0216	0.0134	0.0131	-0.4329	0.1020	0.8408	-0.0039	3.5736	6.3934	0.0129
	[-189.0840]	[55.5097]	[41.9610]	[88.4175]	[102.5569]	[-26.4710]	[2.4606]	[44.8443]	[-16.4765]	[15.1478]	[19.6733]	

Panel B: fast sells

	Const	FastBuys(-1)	FastSells(-1)	SlowBuys(-1)	SlowSells(-1)	netQuotes	Returns	Quotes	Spreads	BuyLayer	SellLayer	R^2
Q1	-0.0375	0.0054	0.0238	0.0048	0.0043	0.1792	0.0688	0.2640	-0.0005	0.9456	0.8175	0.0014
	[-29.2890]	[8.4017]	[17.5911]	[27.5633]	[24.9574]	[6.8977]	[3.5477]	[15.7224]	[-64.3343]	[15.1931]	[15.4807]	
Q2	-0.1235	0.0082	0.0371	0.0073	0.0072	0.3974	0.0457	0.5612	-0.0008	1.5863	1.7615	0.0035
	[-51.3785]	[12.3525]	[12.5901]	[41.2489]	[36.3690]	[10.8332]	[11.1564]	[25.6363]	[-57.5591]	[13.2468]	[19.8712]	
Q3	-0.2632	0.0149	0.0485	0.0083	0.0087	0.3489	0.3125	0.6776	-0.0029	2.4582	2.6580	0.0053
	[-70.9745]	[13.5047]	[19.1862]	[40.7578]	[41.6991]	[9.0736]	[5.1187]	[27.8695]	[-65.2253]	[13.9033]	[18.5051]	
Q4	-0.5841	0.0150	0.0475	0.0105	0.0111	0.3183	0.0166	0.5908	-0.0004	3.5145	3.3077	0.0069
	[-107.2228]	[22.0000]	[24.8656]	[70.3843]	[55.9258]	[15.0764]	[8.8395]	[29.3701]	[-16.6288]	[20.0465]	[22.0206]	
Q5	-1.8585	0.0222	0.0625	0.0137	0.0144	0.4981	0.3058	0.8684	-0.0043	7.7704	3.8420	0.0135
	[-194.3211]	[39.5878]	[57.3366]	[103.5573]	[93.1975]	[29.2962]	[2.1882]	[44.0292]	[-24.3936]	[22.2335]	[15.3613]	

Notes: Stocks are grouped into 5 ADV based quintiles. For each quintile the table report results from fixed effect panel regressions of the number of fast buy or sell trades within an interval of one second on lagged trades of all types, lagged returns, the lagged bid-offer spread, the lagged gross number of quotes (buy orders plus sell orders), the lagged net number of quotes (buy orders minus sell orders), and lagged indicator variables taking the value unity if buy or sell side spoofing has been identified ($BuySpoof_{i,t-1}$ or $SellSpoof_{i,t-1}$). All coefficients, aside from those on the four lagged trade variables, are multiplied by 1000 for legibility. Heteroskedasticity robust standard errors are reported in brackets.

Table 8: Fast and Slow Trades: pre and post-trade changes in depth: same-side: 5 event horizon

	Pre-trade				Post-trade			
	Fast		Slow		Fast		Slow	
	Best	Top5	Best	Top5	Best	Top5	Best	Top5
Q1	-0.047	0.077	-0.065	0.013	0.182	0.193	0.001	-0.076
	[0.023]	[0.032]	[0.011]	[0.012]	[0.049]	[0.062]	[0.013]	[0.027]
Q2	-0.014	0.026	-0.040	-0.035	0.078	0.078	-0.012	-0.084
	[0.011]	[0.015]	[0.011]	[0.013]	[0.029]	[0.042]	[0.016]	[0.018]
Q3	-0.019	0.004	-0.024	-0.027	0.008	0.001	-0.003	-0.054
	[0.004]	[0.007]	[0.003]	[0.004]	[0.004]	[0.015]	[0.002]	[0.009]
Q4	-0.023	-0.021	-0.030	-0.047	0.015	0.019	-0.001	-0.018
	[0.005]	[0.008]	[0.005]	[0.008]	[0.007]	[0.011]	[0.003]	[0.006]
Q5	-0.016	-0.017	-0.015	-0.017	0.002	0.001	-0.001	-0.003
	[0.004]	[0.004]	[0.003]	[0.003]	[0.001]	[0.002]	[0.001]	[0.001]

Notes: results for 5 ADV-based quintiles of the universe of 289 stocks. For each quintile the table presents the average change in depth (first within stock and then across stocks) prior to trades and then presents average depth changes after trades. Depth changes here are for the 'same side' that the trade hits i.e. on the limit sell (buy) side for aggressive buy (sell) trades. Depth changes are computed over 5 order events prior to the trade (from $t-7$ to $t-2$) or 5 order events afterwards (from $t+1$ to $t+6$) at the best and also summed across the first five levels of the order book, respectively. Depth is measured in percentage points of average daily volume traded. Means are presented and below these in brackets are cross sectional standard errors (equal to the standard deviation of the cross-stock averages divided by the square root of the number of stocks in each subgroup).

Table 9: Fast and Slow Trades: pre and post-trade changes in depth: opposite-side: 5 event horizon

	Pre-trade				Post-trade			
	Fast		Slow		Fast		Slow	
	Best	Top5	Best	Top5	Best	Top5	Best	Top5
Q1	0.148	0.277	0.072	0.128	-0.115	-0.283	-0.042	-0.042
	[0.037]	[0.053]	[0.008]	[0.013]	[0.049]	[0.062]	[0.013]	[0.027]
Q2	0.068	0.219	0.048	0.094	-0.054	-0.204	-0.034	-0.042
	[0.012]	[0.025]	[0.004]	[0.009]	[0.029]	[0.042]	[0.016]	[0.018]
Q3	0.029	0.145	0.027	0.056	-0.026	-0.101	-0.021	-0.031
	[0.006]	[0.018]	[0.002]	[0.006]	[0.004]	[0.015]	[0.002]	[0.009]
Q4	0.010	0.077	0.022	0.030	-0.024	-0.064	-0.016	-0.017
	[0.011]	[0.016]	[0.003]	[0.005]	[0.007]	[0.011]	[0.003]	[0.006]
Q5	0.011	0.017	0.009	0.009	-0.011	-0.015	-0.006	-0.005
	[0.003]	[0.003]	[0.002]	[0.002]	[0.001]	[0.002]	[0.001]	[0.001]

Notes: results for 5 ADV-based quintiles of the universe of 289 stocks. For each quintile the table presents the average change in depth (first within stock and then across stocks) prior to trades and then presents average depth changes after trades. Depth changes here are for the 'opposite side' to that which the trade hits i.e. on the limit sell (buy) side for aggressive sell (buy) trades. Depth changes are computed over 5 order events prior to the trade (from $t-7$ to $t-2$) or 5 order events afterwards (from $t+1$ to $t+6$) at the best and also summed across the first five levels of the order book, respectively. Depth is measured in percentage points of average daily volume traded. Means are presented and below these in brackets are cross sectional standard errors (equal to the standard deviation of the cross-stock averages divided by the square root of the number of stocks in each subgroup).

Figure 1: Bid-ask quotes around fast and slow transactions

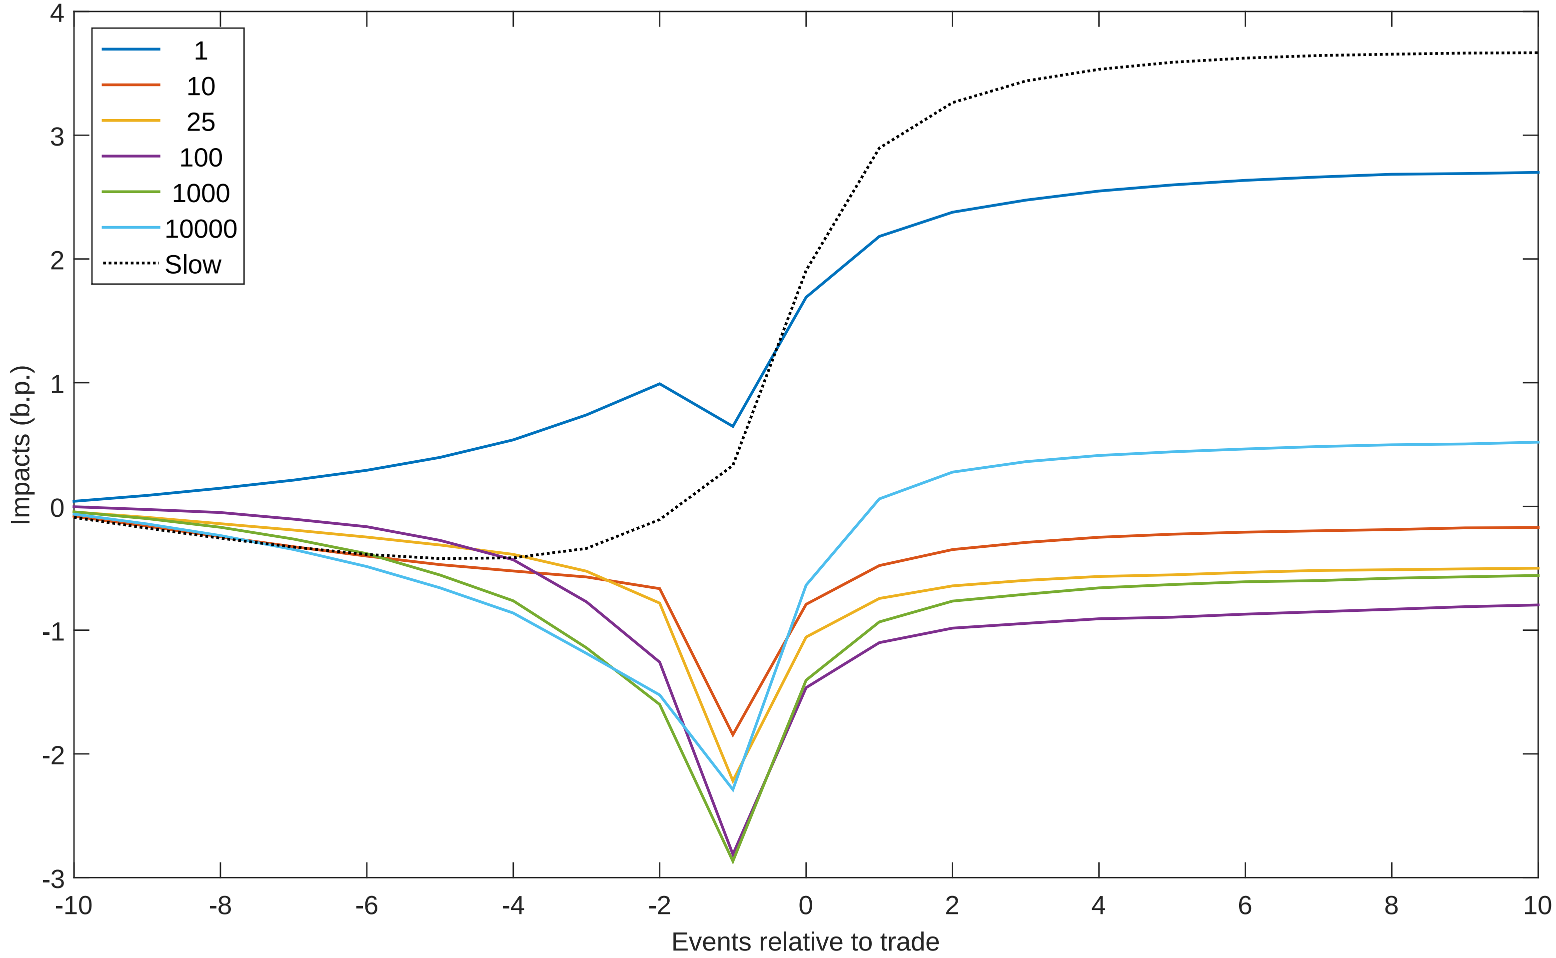

Notes: Time (horizontal axis) is measured in quote time. The mid-price is normalized to zero at quote -15. The quotes prevailing at time 0 are those immediately after the trade has executed.

Figure 2: Intra-day patterns in trading and liquidity measures

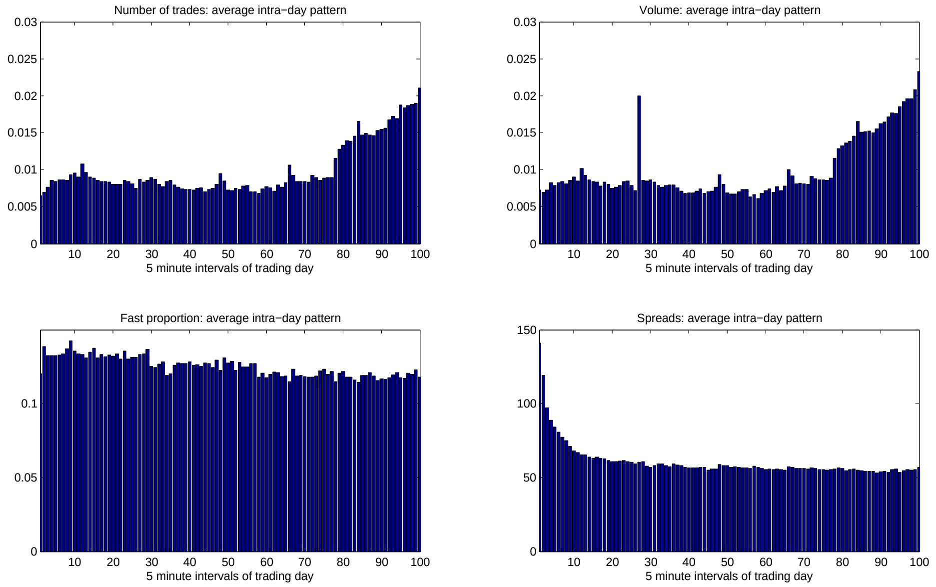

Figure 3: Time from order entry to execution: sorted by average daily traded volume

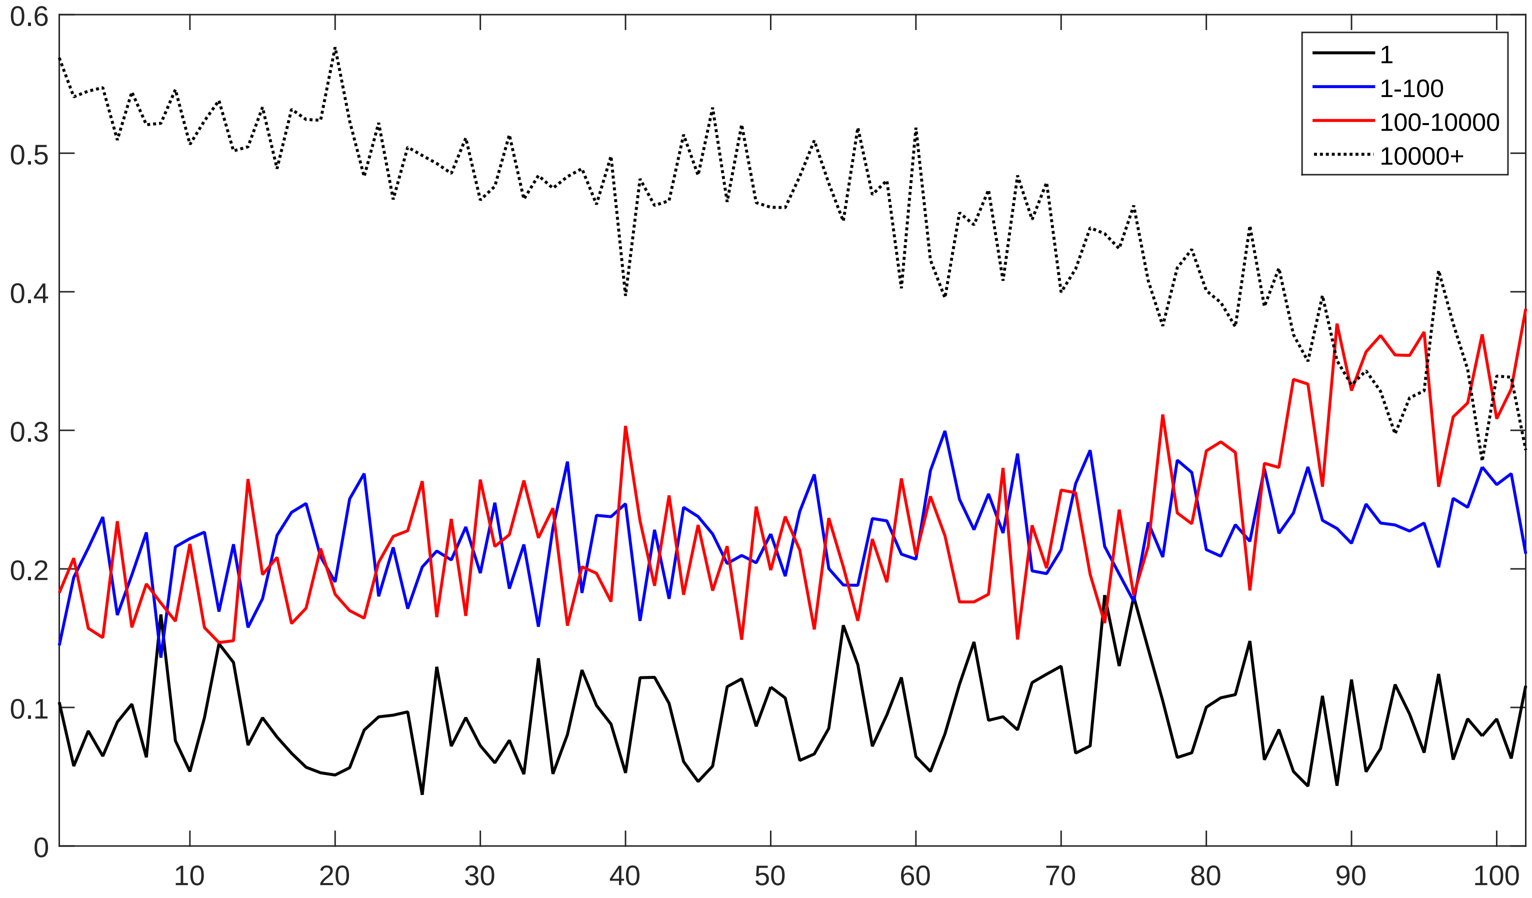

Figure 4: Spoofing event counts by firms, ordered by increasing ADV