Government Office for Science

Foresight

High frequency trading – assessing the impact on market efficiency and integrity

Driver Review DR28

Foresight, Government Office for Science

Contents

Executive summary ... 3
1. Introduction ... 3
2. Definition of HFT ... 4
3. Institutional detail ... 5
3.1. Structural changes at the market level ... 5
3.2. Changes at the exchange level ... 6
4. Literature review ... 7
4.1. Market quality ... 7
4.2. Market efficiency ... 7
4.3. Market integrity ... 10
4.4. Estimates of HFT activity ... 10
5. Data ... 11
6. Market efficiency and market integrity ... 12
6.1. Market quality metrics ... 13
6.2. Market efficiency ... 13
6.3. Price discovery ... 14
6.4. Market integrity ... 15
7. Research design ... 21
7.1. A simultaneous structural equations model for market quality research ... 21
7.2. Stationarity of our estimates and adjustment for Autocorrelation ... 23
8. Results ... 23
8.1. Descriptive statistics ... 23
8.2. Regression results ... 33
9. Robustness test ... 38
10. Conclusion ... 39
References ... 41
Appendix 1: Robustness test regression results ... 43

High frequency trading – assessing the impact on market efficiency and integrity

This report has been prepared for the Foresight Future of Computer Trading in Financial Markets Project series by Michael Aitken, Frederick H. deB. Harris, Tom McInish, Angelo Aspris, Sean Foley through the Capital Markets Cooperative Research Centre (CMCRC)[1]

1 May 2012

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation.

[1] The authors would like to acknowledge the CMCRC development team for assistance in the construction of all metrics and the Securities Industry Research Centre of Asia-Pacific for the provision of data, without which this project would not be possible. The helpful suggestions of two anonymous referees are gratefully appreciated.

Executive summary

High Frequency Trading (HFT) has increased in prominence over the last decade. What was once a novel feature experimentally employed by sell-side institutions is now a commonplace requirement amongst a diverse group of market participants. The technological advancements in financial markets which have been responsible for a significant reduction in market-wide latency and improved data feeds combined with an introduction of new market access models featuring aggressive fee structures have contributed to a steady increase in HFT. Yet, despite the known changes that have facilitated the increase in HFT, very little is still known about such activity and its impact on market quality.

The mandate of securities regulators, who are required to sign off on market design changes, requires that they assure themselves that any such adjustments, of which high frequency trading is an example, must enhance and certainly not detract from the efficiency and/or integrity of markets. Although some work has been done on the impact of market design changes (and HFT in particular) on market efficiency under the guise of market quality, little, if any attention has been directed towards examining how market design changes impact market integrity or the interplay between efficiency and integrity. These issues form the purpose of the research project at hand.

Our findings indicate that HFT now forms a significant proportion of order flow across Europe's major markets. Noticeable changes in market quality metrics representing both efficiency and integrity suggests that HFT may not be the bogey man that many buy-side institutions have made it out to be. Our results show that efficiency has improved as has the integrity of the market in the wake of increases in HFT. The findings in this report should be viewed cautiously given the experimental market quality framework we are employing[2] and the fact that the proxies for HFT, efficiency and integrity were all constructed using the less reliable public rather than private data.

1. Introduction

Despite being in the spotlight for some time, our understanding of high-frequency trading and its implications for market quality are at best, moderately informed. The rapid pace of adoption of high frequency trading, coupled with perceived increases in market disruptions, market volatility and market abuse over the last few years have cast a shadow over the desirability of HFT. In principle, high frequency trading is an innocuous activity designed to allow traders to minimise trading costs (market impact costs), enter into and/or exit market positions with greater flexibility, and execute strategies that profit from speed. Though the practice reflects a natural progression in markets, from the traditional floor-based trading venues to a system of inter-connected markets where best execution is normally determined by way of prices quoted across regulated markets, multilateral trading facilities and alternative trading systems, its existence has been linked with some more sinister outcomes. For example, the "Flash Crash"

[2] See Aitken, M, J and R de Harris, Evidenced-Based Policy Making for Financial Markets: A Fairness and Efficiency Framework for Assessing Market Quality, with Rick Harris, *Journal of Trading*, Vol 6, No3, Summer 2011, pp.69-89.

that hit US equities markets on 6 May 2010, wiping 650 points off the Dow in half an hour, has been in part attributed to practices of high-frequency traders. [3] While regulators make use of formal consultation processes to help them make decisions about whether to authorise a given market design change, the nature of the evidence provided is normally supported by an underwhelming amount of hard evidence linking the change to the overall efficiency and integrity of markets. Addressing these concerns, our work deploys a new market quality framework within which both efficiency and integrity are defined and measured. [4]

The principal question being addressed in this paper is how HFT has impacted the efficiency and integrity of the LSE and NYSE-Euronext Markets in Paris. Using data from 2003-2011 we address this question by deploying a unique 3SLS framework that captures the interrelationship between market efficiency, market integrity, and HFT. Using this framework we are able to show that HFT participation has had a significant impact on both the cost of trading (our proxy for efficiency) and the degree of end of day price dislocation (one of our proxies for market integrity), both of which have improved as HFT participation has increased.

This evidence adds to the regulatory debate on high-frequency trading by showing that market design changes, such as the introduction of HFT, has implications for market efficiency and market integrity. Any assessment of a design change should therefore assess not only the impact on each individual element, but also consider the interplay between the multiple dimensions of market quality. The remainder of this paper is structured as follows. Section 2 explores a wide range of sources from consultation papers to academic pieces to explain what HFT is and have it has developed. Section 3 provides an overview of the institutional detail in the European marketplace. Section 4 examines the academic literature and provides a review of the main findings. Section 5 introduces the data used for the analysis and section 6 describes the various proxies designed to assess market efficiency and integrity. Section 7 describes the research design and section 8 presents and discusses the results. Section 9 reports the results of our robustness analysis.

2. Definition of HFT

Though modest attempts have been made by regulators, academics, and practitioners to define HFT, most statements that appear in circulation seem only to highlight the fact that it is a tool employed by professional traders. [5] More recent and earnest efforts to understand and define HFT, have however, recognised the multifaceted nature of the activity which can be used to produce more accurate and thoughtful explanations.

Confused in many open debates is the relationship between HFT and another contentious form of trading activity: algorithmic trading. The two, though possessing some similarities, are not synonyms. Rather, HFT is generally viewed as a subset of algorithmic trading (see Brogaard, 2010) where trading decisions are normally pre-designed and submissions are automated and executed without human intervention. Whereas algorithmic trading is fundamentally employed

[3] Unsurprisingly this was followed by a significant and immediate rebound (Kirilenko, Kyle, Samadi and Tuzun (2010)).

[4] See footnote 1.

[5] For a comprehensive list of definitions see Gomber, Arndt, Lutat and Uhle (2011)

in agency trading to achieve particular outcomes such as stealthily capturing liquidity, engaging in block trading in a manner that minimises information leakage, or simply minimising implementation shortfall, HFT is specific to proprietary traders. [6] It is characterised by rapid order submissions and cancellations, which are normally accommodated by exchanges offering co-location services (thereby facilitating low latency). Furthermore, its participants hold little to no over-night positions. This minimum inventory position characteristic adds some insight to the core business model used by those that engage in HFT. HFT traders realise profits from small deviations in prices across hundreds if not thousands of transactions throughout the day. It is therefore, a low margin, high turnover trading practice.

Similar to algorithmic trading, HFT is characterised by a broad range of active strategies employed by a diverse group of trading participants. These participants range from proprietary market making firms to quantitative hedge funds and their practices include a range of activities such as pseudo market-making and statistical arbitrage ("stat-arb") trading. Though many of these trading approaches derive from strategies that have always existed in markets, the speed with which they are able to be employed as well as the cost of doing so are what has changed significantly. [7] With these changes delivering significant competitive advantages to HFT participants, the dramatic growth in HFT over such a short time horizon is comprehensible.

3. Institutional detail

Regulatory framework changes have provided both the impetus for and restrictions on this nascent form of trading practice. As an increasing number of firms enter this space, seeking to carve out niches for themselves, regulators will be required to keep pace with the movement by trying to understand how HFT impacts on market fairness and efficiency. This section provides a broad examination of the major structural changes at both the market and exchange level that have facilitated the expansion of HFT participation.

3.1. Structural changes at the market level

Following a similar transformation in U.S. financial markets some years prior, the implementation of the Markets in Financial Instruments Directive (MiFID) in 2007 transformed the market landscape in Europe. Though the changes were not designed to promote HFT, the variation in rules regarding best execution, market transparency, and market organisation resulted in this outcome. According to a statement from the Committee of European Securities Regulator (CESR):

> "...one of the primary objectives of the MiFID was to introduce competition among trading venues and thereby reduce transaction costs for the benefit of both the final investors and issuers. These results were meant to be achieved, among others, by removing the concentration rule and, at the same time, introducing specific pre- and post-trade

[6] Proprietary traders (or "prop traders") utilise their own capital for trading activities to create profit.

[7] Whilst the nature of trading strategies employed has not significantly changed over time, the mix of participants organising and executing HFT strategies has fundamentally evolved over this period.

transparency for listed shares coupled with a detailed best execution regime for investment firms which execute clients' orders. [8]"

MiFID specifically introduced four key changes with clear implications for HFT. These changes included:

1. The abolishment of the 'concentration rule';
2. The enacting of 'best execution' obligations;
3. The promotion of greater pre- and post-trade transparency rules;
4. The introduction of 'passporting'.

The abolishment of the concentration rule, under which, in some European countries, required all trades to go through a regulated exchange, ensured that there would be greater competition amongst trading venues in Europe. Provisions in MiFID to allow competing structures (multilateral trade facilities (MTF's) and systematic internalises (SI's)) saw traditional incumbent exchanges being challenged by more aggressive entrants who were supported by the introduction of a principles based best execution regime. [9] This change imposed a requirement on the incumbent exchange to compete on cost. Market operators, looking to attract order flow would have to offer alternative fee structures that were typically asymmetric in nature. Under such conditions, market participants that provided liquidity enjoyed a lower fee (and in some circumstances a rebate) relative to those participants that demanded liquidity from the market. [10] The outcome of such a change was expected to deliver a market environment that was characterised by lower trading costs and evidence provided by Gomber et al. (2011), suggests that the delivery of this fragmented, yet evidently more harmonised marketplace met this expectation. Given the acute sensitivity of HFT to the cost of trading, regulatory changes which have allowed both implicit and explicit trading costs to fall, have further encouraged more involvement by HFT participants.

In the aftermath of these initial changes, the European Commission's second version of the MiFID directive (MiFID II) imposes a line of future challenges to the practice of HFT. In a bid to safeguard against potential systemic risks, the directive warned that there was now a greater potential for "overloading of trading venues systems", "duplicative or erroneous orders", and "overreacting to other market events". [11] Though these risks are still the subject of much contention, proposals to require entities who engage in HFT to employ greater risk mitigation practices may produce a range of yet unintended consequences. [12]

3.2. Changes at the exchange level

Whilst changes at the market level have been partly responsible for driving the increase in HFT participation, at the exchange level, a swathe of initiatives have also supported this dramatic

---

[8] Committee of European Securities Regulator (CESR) – 24th April 2006.

[9] This was subsequently amended in MiFID II bringing it in line with the U.S. regulatory regime.

[10] These market operators would be able to offset some of the cost of the rebate from 'tape-revenue'.

[11] MiFID II – Committee of European Securities Regulator (CESR)

[12] HFT is also been scrutinised greatly at the national level, with a recent proposal in France requiring that a 0.1% tax on firms that cancel more than 50% per day of their automated orders be implemented.

rise. For example, in June 2007, the London Stock Exchange (LSE) introduced a system called *TradElect*, which promised to deliver an average 10 millisecond turnaround time (latency) from placing an order to final confirmation. [13] This system was subsequently replaced by *MillenniumIT*, which went live in 2011 delivering average latency of less than 120 microseconds, making LSE one of the fastest trading venues in the world. Unlike the initial MiFID directives, the increased investment in market infrastructure was geared towards capturing a greater share of the HFT order flow. For example, in 2007, the LSE also introduced Member Authorised Connection (MAC) – a direct, low latency connectivity solution for non-member firms. This solution allowed non-member firms to access the Exchange's centrally cleared order driven Trading Services under an authorising member firm's trading codes. Despite adopting and abandoning a maker-taker fee structure in 2008, the LSE, responding to pressures from competing trading platforms like Bats Europe, introduced an amended tiered pricing facility based on trade turnover. In light of active efforts by exchanges to push low latency solutions, accompanied by lower trading costs, and greater market access solutions, HFT participation will continue to feature prominently across European markets.

4. Literature review

Despite the widespread use of algorithmic and high frequency trading in today's modern capital markets, academic research in this area is only just beginning to scratch the surface. Research in this area suffers from a paucity of data, primarily due to the lack of market-wide HFT indicators that would greatly facilitate such research. Whilst firms engaging in HFT activity could provide their own internal data, the sensitivity of the trading strategies employed by such proprietary algorithms has yielded only one piece of research to date. The studies that do exist use a variety of proxies to determine the extent of HFT participation and find that it occupies a significant position in the market.

4.1. Market quality

While most academics associate market quality with efficiency, according to the IOSCO 2003 statement on the Objectives and Principles of Securities Regulation, market quality should include a focus not only on market efficiency but also market fairness/integrity. Any assessment of market quality that does not address both aspects is therefore only dealing with one side of the debate.

4.2. Market efficiency

A small but increasing number of theoretical and empirical studies have contributed to the pool of evidence on the impact of HFT on market efficiency. Of these papers, the focus has largely been on the cost of trading rather than the price discovery element of efficiency. The foremost theoretical study in this area, by Cvitanic and Kirilenko (2010) suggests that the introduction of HFT participants will reduce the average trade value and move the distribution of prices closer to the mean, resulting in reduced volatility. Their model incorporates an electronic limit order book in which both high-frequency (HFT) and low-frequency (human) traders participate. A key assumption of their model is that the HFT participants act as uninformed market-makers. Their findings also indicate that the time between trades and the volume of each trade decreases in proportion to the participation of human traders. That is, more frequent human trades result in lower average trade volume and lower duration between trades.

[13] Or equivalently accommodating 3000 orders per second.

Jovanovic and Menkveld (2010) construct a theoretical model that provides conflicting findings to those reported in Cvitanic and Kirilenko (2010). In their model, Jovanovic and Menkveld (2010) do not assume a "classical" Kyle (1985) style market maker model, in which the "middlemen", HFT participants, are uninformed. Rather, they assume that HFT participants are both faster *and* more informed than their counterparts. With this assumption relaxed, their findings for market efficiency are mixed. Their theoretical model suggests that HFT participants may cause welfare gains by increasing efficiency through reduced spreads. This will be achieved by updating their information set faster than a traditional market maker, thereby increasing their ability to avoid adverse selection. This ability to constantly update their information set allows HFT participants to charge a smaller fee for the provision of liquidity, reducing spreads, whilst at the same time avoiding stale orders entering the market, resulting in increased efficiency. On the other hand, Jovanovic and Menkveld (2010) warn that if no adverse selection problem exists to begin with, the ability of HFT participants to update their information set faster than human traders may in fact cause an adverse selection problem, reducing the willingness of human participants to enter the market for fear of trading with a better informed HFT participant. The authors theorise that such a situation could widen spreads and thereby reduce efficiency. In their empirical analysis, they consider 77 trading days of Dutch index stocks between 2007 and 2008 to study the effect of the introduction of Chi-X, which is used as an instrument for the escalation of HFT participation. Their evidence suggests that the HFT participants in their sample are in fact better informed about news than the average trader. In calibrating their theoretical model using the available empirical data they determine the addition of HFT participants has caused a marginal increase in welfare over the examined period.

Some potential negative effects of HFT participants are also outlined in a recent study by McInish and Upson (2012). The authors develop a theoretical framework modelled on US equity markets, explicitly examining the SEC's Flicker Quote Exception rule. This rule allows inter-market trade throughs to occur as long as the new price has been displayed for less than one second. The theoretical component of the paper demonstrates how this rule allows HFT participants to profit from their knowledge of the "true" current state of the market against "slower" human participants. This informational advantage allows the HFT participants to "pick off" orders from slow liquidity demanders in the sub-second environment at prices inferior to the national best bid and offer (NNBO). Whilst this paper explicitly deals with the US market, such situations would also be possible in any market where trade through prohibitions do not exist.

Despite these largely theoretical negative market quality findings, the majority of papers that deal with HFT empirically find a predominantly positive overall impact. The predictions made by Cvitanic and Kirilenko (2010) of reduced volatility, trade value and volume are empirically supported by the work of Jarnecic and Snape (2010). This study uses a proprietary dataset from the London Stock Exchange (LSE) containing unique identifiers of the type of participant behind each order. These participants are separated into high frequency traders, investment banks, large institutions, small institutions, retail brokers and market makers. This separation allows the authors to describe the characteristics of liquidity supply and demand by HFT participants. Their results indicate that HFT participants both contribute and demand liquidity in almost even proportions, and that their activity is more likely to dampen than increase volatility.

The question as to how HFT impacts on informational asymmetry costs, modelled theoretically and empirically by Jovanovic and Menkveld (2010), is also empirically addressed by Hendershott, Jones and Menkveld (2011). In this paper, the authors use the volume of message traffic, normalised by the number of trades, as a proxy for the level of HFT participation. Their five year sample period straddles the staggered introduction of the automation of quote dissemination on the NYSE in 2003. Their findings indicate that algorithmic trading, of which HFT is a subset, significantly reduces both quoted and effective spreads (market efficiency). They find that this narrowing of the spreads is the result of a decline in adverse selection and an increase in the price discovery associated with the trades. Consistent with the results of McInish and Upson (2012), Hendershott, Jones and Menkveld (2011) find that the introduction of algorithmic trading increases the profits to liquidity providers, through increased realised spreads.

Evidence gleaned directly from HFT firms tells a story similar to that found by Jarnecic and Snape (2010) and Hendershott, Jones and Menkveld (2011). Brogaard (2010) uses a dataset provided by NASDAQ that directly identifies 26 HFT traders over periods during 2008, 2009 and 2010. Brogaard addresses the concern that HFT participants flee the market during times of heightened volatility. Each trading day in his sample period is divided into 15 minute segments, allowing him to separately analyse periods of both high and low volatility. He finds that when prices fluctuate more than normal, HFT participants supply more liquidity and demand less liquidity, than on average. Brogaard (2010) uses this as evidence that on average HFT participants are unlikely to exacerbate volatility, consistent with the findings of Jarnecic and Snape (2010).

Brogaard (2010) also uses this dataset to analyse the price discovery attributable to HFT and non-HFT participants by utilising both the variance decomposition metric constructed by Hasbrouck (1991) and the information share metric introduced by Hasbrouck (1995). Results using these metrics suggest that the price discovery contribution of HFT participants is greater than that of the non-HFT participants. This implies that HFT participants contribute to the efficiency of the market. Brogaard (2010) furthermore conducts an analysis of the liquidity provided by HFT participants by looking at the depth and time spent by each participant type at the NBBO. His findings indicate that whilst HFT participants spend more time at the NBBO, they provide less depth than their non-HFT counterparts. They additionally find that HFT participants are better able to avoid trading with insiders than are non-HFT participants, consistent with the theoretical findings of Jovanovic and Menkveld (2010).

One of the most comprehensive datasets has come from trading in German stocks on Xetra. Due to a rebate scheme applied to algorithmic traders, a flag is applied to every trade identifying if it arises from a human or algorithmic source. Groth (2011) uses four trading days during 2007 to analyse the conduct of HFT participants during times of high and low volatility, by splitting the day into 5 minute intervals and identifying the level of volatility in each. His findings indicate that HFT participants do not change their trading behaviour conditional on the volatility in the market and that they are as active in periods of high liquidity as they are in periods of low liquidity. Groth (2011) also finds that there is no evidence of market withdrawal by HFT participants during periods of increased volatility.

One of the more unique proxies for HFT used in the literature is described by Hasbrouck and Saar (2010). The authors use trade and quote data in the millisecond environment to identify strategic runs of trades, observing interactions between traders separated by as little as 3-5 milliseconds. These strategic runs are used to create a proxy for the level of HFT participation. Hasbrouck and Saar (2010) analyse the trading activity on the NASDAQ for one month each during both 2007 and 2008, which unlike studies such as Groth (2011) and Jarnecic and Snape (2010) allows them to analyse HFT participation during a period of high market stress. Similar to other HFT studies, Hasbrouck and Saar (2010) divide the day into 10 minute intervals, however they uniquely apply a two stage simultaneous equation for each market quality metric, allowing for a potentially endogenous relationship between HFT participation and their market quality indicators of volatility, depth and spreads. Their findings indicate that increased HFT participation increases quoted depth; reduces quoted spreads and reduces volatility, even in times of market stress.

Whilst the proxies used for HFT activity differ significantly throughout, all but one of these empirical studies documents positive impacts on the market efficiency aspect of market quality. Specifically, the evidence provided suggests that increased levels of HFT participation result in narrower spreads, greater price discovery and reduced volatility. These findings must be considered preliminary at best due to the absence of hard data on the true level of HFT activity. Striking is the absence of any studies to date that analyse the impact on market integrity (a key component of market quality as described earlier), despite concerns from market participants regarding the potential predatory behaviour of some of the algorithms used by HFT participants.

4.3. Market integrity

Despite a raft of statements from market participants suggesting that HFT may be having a detrimental impact on the integrity of the marketplace we are not aware of any papers that address this issue. To begin addressing the issue we first define market fairness/integrity as the extent to which market participants engage in prohibited trading behaviours. This requires us to measure insider trading, market manipulation and broker-client conflict (e.g. front-running). None of these are directly observable forcing us to identify proxies such as information leakage (for insider trading) and the dislocation of the end of day price at month/quarter end and on option expiry dates (for market manipulation)[14]. These proxies are explained in more detail later in the paper.

4.4. Estimates of HFT activity

The empirical evidence regarding the size of HFT participation is relatively sparse and relies primarily on noisy proxies. Perhaps as a result, the estimates that have been provided in the literature vary widely. On the LSE, Jarnecic and Snape (2010) find that from their 2009 sample, 40-64% of trades included an HFT participant on at least one side. In their 2010 response to the Committee of European Securities Regulators (CESR) call for evidence, the LSE identified that during 2010 their internal estimates of HFT participation varied between 32-33% of total UK equities trading. In a similar submission, NYSE Euronext calculated that in the overall European market there was a 5% market share for HFT participants in the first quarter of 2007, increasing to 23% of total traded value in the first quarter of 2010.

Brogaard's (2010) analysis, which focuses on U.S equities, finds that 60-80% of all NASDAQ trades involve an HFT participant as either a liquidity provider or demander. Ito and Lyden (2012) construct an undisclosed measure of HFT participation for the largest 15 stocks traded on NASDAQ, NYSE and BATS in the US and show that HFT participation in one side of trades between 87-92% of the time.

[14] Because we use publically available data note that proxies for broker-client conflict (i.e. front-running) cannot be estimated at all. This occurs because most markets either don't have rules requiring brokers to identify when they trade as principal or agent, or as in London where they do have such rules, the rules are easily circumvented by brokers marking all trades as principal trades and only differentiating at the settlement level. Further, rough proxies for front-running are only available where broker/trader codes are provided and that does not generally occur with publically available data.

As mentioned, these estimates of HFT participation in the market vary significantly. One thing that is implicitly agreed upon by all researchers addressing the question of HFT participation is that these participants are able to extract value from the market. There are a few papers that attempt to empirically estimate these profits, with the results as diverse as the estimates on trading activity.

The first paper to address this question is Kearns, Kulesza and Nevmyvaka (2010). The authors analyse a broad cross section of trades on all US markets and construct an optimistic upper bound for the maximum profits that could be attributed to HFT participants during the year 2008, without utilising any particular HFT proxy. They find that this upper bound to be approximately US$21 billion. This estimate is roughly 10 times larger than the annual $2.8 billion computed by Brogaard (2010) using his dataset of actual HFT trades. McInish and Upson (2012) conduct an empirical analysis of inter-market arbitrage using 2008 NASDAQ data. This arbitrage is made possible within the sub-second environment by the Flicker Quote Exception rule. They estimate that over $233 million per year is transferred from slow retail traders to fast HFT participants in the US due to this phenomenon.

Several studies, including Brogaard (2010), Jarnecic and Snape (2010) and Ito and Lyden (2012) find that HFT participants are more active in larger stocks than smaller stocks, and are comparatively more active towards the end of the day. These results are found to be indicative of the market making nature of the HFT participants, seeking to close the day with zero inventory positions, if possible.

To date, very few event studies have been conducted in relation to HFT participation. Kirilenko, Kyle, Samadi and Tuzun (2011) provide the most recognised such study dealing with HFT, utilising audit trail data for the E-mini S&P 500 futures contracts on the day of the "Flash Crash", May 6^{th}, 2010. This data is used to identify HFT participants. Their objective designation relies on trade frequency and size and the authors use this to determine that 16 accounts out of a total of 15,422 belong to HFT participants on that day. By analysing a variety of metrics including holding periods, inventories and trade directions, they find that HFT participants, after providing some initial liquidity to fundamental sellers, contributed to the significant selling pressure that precipitated the flash crash. Whilst Kirilenko et al. (2011) do not go so far as to blame this incident on HFT participants, they do determine that their presence in the market exacerbated the volatility present during this period of extreme market stress.

This evidence stands in contrast to an experiment using actual HFT data conducted by Brogaard (2010). In this study, the author examines the supply of liquidity by HFT participants during the four days in September 2008 in which the news regarding the collapse of Lehman brothers became public. Brogaard (2010) finds that HFT participants did not significantly increase their demand for liquidity, but did significantly increase their liquidity supply. This analysis, and an examination of HFT participation following earnings announcements with similar findings, leads Brogaard (2010) to conclude that HFT participants do not remove liquidity from the market, even in times of severe market stress.

5. Data

The data used in this study encompasses all listed securities on the LSE and Euronext-Paris between 2001 and 2011. The trade and quote data used for the analysis is obtained from Thompson Reuters Tick History (TRTH) which is provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA).

In relation to the fields obtained, we acquire information regarding bid and ask quotes (time stamped to the nearest millisecond) for up to ten levels of the order book. This is captured for every security listed on the LSE and Euronext Paris. We access intra-day information related to market prices, volume, number of level 1 orders/trades, as well as all changes to the level 10 orderbook.

To construct our measure of market integrity (informational leakage) we acquire company news announcements from the Reuters News Announcement Feed (RNA). These announcements only became available after 2003 so we have restricted our analysis to the period 2003 – 2011, inclusive.

In order to determine the number of order cancellations (to be used to proxy the level of HFT) it is necessary to match trade executions against the limit order book immediately prior to trades. To match these trade executions, we follow the methodology employed in Hasbrouck and Saar (2010) in assuming that the dominant economic event is the arrival of a marketable limit order. Similar to their dataset, when an incoming marketable order executes against more than one standing limit order multiple messages are generated for each standing limit order. We combine each marketable order arrival, defined as multiple order executions with the same millisecond time stamp, in the same direction that are unbroken by any non-execution message, as a single marketable order.

The cancellations data used in this study are generated by analysing trade and level 10 quote data. Any decrease in quoted depth is caused by either an execution against an incoming market order or a cancellation or amendment request. We match all trades and quotes by time and price to account for all quoted depth reductions due to trade executions. We then analyse the change in quoted depth for the remaining quote updates and define each into one of the following three categories:

1. If the new quote update only increases depth at any of the 10 levels from the previous quote then it is considered to be due to the addition of a limit order.
2. If the quote update contains additions to depth at some levels and reductions in depth at others, this is considered to be an amendment.
3. If the quote update contains only reductions in depth at any of the 10 levels and is not associated with a corresponding trade it is considered to be due to a cancellation.

These cancellations are then compared to the number of trades to arrive at the Cancel to Trade ratio which we use to identify the rate of HFT participation. This cancellation ratio has been used in previous empirical works such as Hendershott, Jones and Menkveld (2011) to identify the level of HFT participation, and is found to be highly correlated with explicit measures of HFT due to the high rate of order submissions and cancellations necessary to conduct a HFT strategy.

6. Market efficiency and market integrity

The following section describes the metrics that will be used in our empirical analysis. This section also describes the rationale for the proxies adopted as well as the methods used to construct them.

6.1. Market quality metrics

The academic literature to date has largely focussed on the impact of HFT on elements of market quality that relates to efficiency; namely spreads, depth, volatility and to a lesser extent price discovery. Whilst these elements are fundamental to the operation of efficient markets, the mandate of national regulatory agencies and exchanges alike is to promote both efficient and fair marketplaces. This study seeks to address the impact of HFT participation on both of these market quality outcomes. The evidence thus far indicates that increasing HFT participation has resulted in increased efficiency. Ideally, this increase in efficiency should occur without harming the integrity of the market. Indeed, one of the main criticisms raised by market participants against HFT is the impact of HFT trading on the fairness and transparency of the marketplace. We utilise a variety of integrity metrics to address this question, with each of the metrics detailed below.

6.2. Market efficiency

Market efficiency is defined as minimising transaction costs whilst also maximising price discovery: the lower the cost of trading and the more quickly information is impounded into price, the more efficient the market. It follows logically that measures of transaction costs and price discovery are useful proxies for market efficiency. The first is estimated by measuring elements of transaction costs in a marketplace such as quoted, effective and realised spreads. Alternative proxies are sometimes used, such as market depth (the revealed order book), market depth at the best bid-offer or latent liquidity (liquidity which is layered up and down the book awaiting the price to get to an appropriate level). Because full order book data is not available to us at this time, these alternative measures can neither be constructed nor compared.

6.2.1 Effective Spreads

Transaction costs are measured by the volume weighted effective spread calculated on a per-trade basis and averaged across the month. The percentage spread is the same statistic, but expressed as a percentage of the instrument's price. We estimate round-trip trading costs as relative effective spreads, measured as the trade price minus the midpoint of the bid-ask spread immediately prior to trades, multiplied by two to reflect both entry and exit trade as in Venkataraman (2001)[15]. We transform this number into a percentage of the share price to arrive at our final statistic.

Relative Effective Spread = 200 x D * ((Price – Mid)/Mid)

- **D:** is the direction of the trade. A value of 1 is given for a buyer initiated trade and a value of -1 for a seller initiated.
- **Price:** the trade price for stock *i* at time *t.*
- **Mid:** is the midpoint price of the ask and the bid for stock *i* at time *t.*

[15] Venkataraman, K., 2001, Automated versus Floor Trading: An analysis of execution costs on the Paris and New York Exchanges, *Journal of Finance*, Vol. 56, No. 4, pg. 1445-1885.

This metric is constructed for each trade on each stock traded on the respective market for each day in the month separately. The average daily spread for each stock is then constructed as the volume-weighted average of the effective spread on each trade, for each stock. This daily-stock effective spread is then converted to a daily-market-wide effective spread by equally weighting the daily-stock spread experienced by each of the stocks traded on that day. In order to reach the monthly-market wide effective spread utilised in this report, the average of each daily-market-wide effective spread is computed for every trading day in the month.

To determine the direction of trades, we rely on the Lee-Ready (1991) algorithm. Specifically, all trades that occur above or equal to the prevailing ask price are classified as purchases and all trades that occur below or equal to the bid prices are classified as sales.

Realised spreads measure the benefits associated with supplying liquidity to a market rather than demanding it. They are estimated by assuming that one has a choice over whether to demand or supply liquidity and one might be expected to be compensated for supplying liquidity. Relative realised spreads are measured in the following way:

Relative Realised Spread = 200 x D*((Price – V)/Mid)

- **D:** is the direction of the trade. A value of 1 is given for a buyer initiated trade and a value of -1 for a seller initiated.
- **Price:** the trade price for stock *i* at time *t.*
- **V:** is the midpoint price 5 minutes after time *t* for stock *i*. [NB: there is debate in the academic community as to the length of time in the future to consider.
- **Mid:** is the midpoint price of the ask and the bid for stock *i* at time *t.*

Whilst the effective spread is positive by construction, the realised spread may be either positive or negative. As the realised spread is indicative of the profits attributable to market markers, a finding of negative realised spreads indicates a *loss* on the part of liquidity providers. Menkveld (2012) explains how this situation may occur when the rebates provided to liquidity providers are of a larger magnitude than the average positioning loss.

6.3. Price discovery

In this version of our report we do not compute measures of price discovery due to the unavailability of specific market data. Within the next few months two new proxies for price discovery will be calculated allowing us to extend the analysis of HFT participation to this element of market efficiency[16].

[16] The two variables will enable comparison of markets that trade homogenous and non-homogenous securities. The former will deploy techniques based on Gonzalo Granger (1995) and Hasbrouck (1995) and measure information share and common factor share across markets. The later will estimate volatility pre and post all information announcements in a given market measuring attributes like the total area under the curve (to reflect the general level of volatility in a market) and the half-life of volatility in order to reflect the time it takes for volatility to return to normal after an information shock.

6.4. Market integrity

We define market integrity as the extent to which market participants engage in prohibited trading behaviours. Chief among these prohibited trading behaviours are insider trading, market manipulation and broker-client conflict (e.g. front-running). While none of these concepts are directly measurable, there are proxies which are highly correlated that allow us to generate an informed opinion of its existence. They include:

1. **Information leakage**: This provides an indirect measure of the level of insider trading.
2. **Dislocation of the end of day price** (particularly around end of month and option expiry dates): This provides an indirect assessment of the degree of market manipulation.
3. **Front-running**: Unfortunately, this cannot be even proxied with the publically available data available to this research.

6.4.1 Information leakage

We consider information leakage as a proxy for market integrity. Information leakage relates to the dissemination of private price sensitive information ahead of a public transmission of the news. The theoretical and empirical findings of Jovanovic and Menkveld (2010) suggest that it is possible for HFT participants to be better informed than human participants. There are several reasons that this could occur. It is possible that human informed traders are trading on their information, resulting in an order imbalance that the HFT participants are then identifying and trading on. Whilst this type of conduct is not directly attributable to the HFT participants, their role in exaggerating the trades of informed participants may exacerbate any information leakage, leading to larger pre-announcement run-ups. Kirilenko et al. (2010) find evidence consistent with this theory around the flash crash of 2010.

A second reason HFT participation may exaggerate the appearance of information leakage is related to the ability of mechanical processes to almost instantaneously incorporate changes to public information. With information providers such as Reuters and Bloomberg now selling RSS feeds in a format that is friendly to mechanical analysis, HFT participants may simply be able to impound the impacts of new information quicker than non-HFT participants. In the worst case, HFT participants' informational advantage could be based on their operators illegal inside information. The true level of insider trading is not directly measurable. We can, however, use the concept of information leakage to reasonably estimate it. Both concepts seek to measure the extent to which price sensitive information ends up in the hands of a privileged few rather than the market as a whole, but only insider trading is illegal. A proxy for information leakage gives us an idea of the extent of the practice and we consider it an upper bound for the incidence of insider trading in a market.

We build proxies for leakage by measuring the extent of unusual price and volume behaviour prior to price sensitive announcements. The leakage metric is calculated using abnormal price returns relative to a market index.

$$Abnormal\ Return_{\downarrow}(i_{\downarrow}t\ ) = Return_{\downarrow}(i_{\downarrow}t\ ) -\ \beta_{\downarrow}i\ (Return_{\downarrow}(m_{\downarrow}t\ )\ )$$

Where $Return_i$ is the daily return on stock *i*, β_i is the stock-specific correlation with the market return constructed over the benchmarking period and $Return_m$ is the daily market return. This measure is constructed to detect prohibited trading behaviour in equities markets. It uses both a market model and clean event window (similar to the FSA's occasional paper 25 Measuring Market Cleanliness)[17]. The leakage analysis is based around price sensitive company announcements, where an announcement is regarded as price-sensitive when the return of the underlying security between t-6 and t+2 is more than three standard deviations (σ_i) away from the average seven-day abnormal return $\overline{(seven\ day\ return_i)}$ constructed using a bootstrap procedure during a 250 trading day benchmarking period ending at t-10. The bootstrap procedure computes 1,000 9-day returns from the benchmark returns period, with replacement.

$$\sum_{t=-6}^{t=+2} Abnormal\ Return_{i_t} > \overline{seven\ day\ return_i} + 3 * \sigma_i$$

For an announcement to be included in the sample, it must also occur within a 'clean' event window. We define clean event windows as the lack of any other stock-specific announcements within the six day pre-event window. Only announcements from clean events are investigated in order to ensure that any price reaction is most likely to be associated with one single announcement. Where the company has more than one announcement within a clean event day, only the first announcement is considered. A case of suspected leakage is defined as an event where there is an abnormal price movement on one or more days in the t-6 to t-1 pre-announcement period in the same direction as the overall announcement return. An example is given below for a positive announcement.

$$\overline{AR}_i = \frac{1}{250} \sum_{t=-10}^{-260} Abnormal\ Return_{i_t}$$

$$Information\ Leakage\ if\ any\ Abnormal\ Return_{i_t} > \overline{AR}_i + 3 * \sigma_i$$

$$Where\ day\ t\ occurs\ between\ event\ day\ -6\ and\ -1$$

$$and\ there\ are\ no\ days\ in\ which\ Abnormal\ Return_{i_t} < \overline{AR}_i - 3 * \sigma_i$$

$$for\ any\ day\ in\ the\ pre-event\ window\ -6\ to\ -1$$

A daily price movement is considered abnormal if it is at least 3 standard deviations ($3 * \sigma_i$) away from the mean abnormal return during the 250 day benchmarking period ending at t-10. In the case of positive price sensitive announcements, at least one pre-event return must exceed the threshold where none of the other pre-announcement days exhibit abnormal returns lower than three standard deviations below the mean. Figures 1 and 2 illustrate contrasting scenarios around the release of a price-sensitive information announcement.

[17] Dubow and Monteiro, 2006.

Figure 1. No informational leakage

Figure 2. Informational leakage

Specifically, Figure 1 shows a significant change in price and volume dynamics occurring after the announcement shown at the top of the screen. Figure 2, however, shows unusual price and volume movements preceding the announcement.

By examining the ratio of suspected information leakages to clean event windows we are able to construct a measure of information leakage that increases as a market becomes more "leaky". The measure is computed as follows:

$$\%\ Information\ Leakage = \frac{\#\ of\ information\ leakages}{\#\ Clean\ event\ windows}$$

Figure 3 presents the number of information leakages for Euronext and LSE respectively. This metric has been constructed on a monthly basis and describes the number of suspected cases of information leakage per market, derived from a sub-set of price sensitive price announcements.[18] The figure presents similar patterns for the LSE and Euronext Paris, with the former displaying elevated signs of informational leakage between 2007 and 2009.

Figure 3. Suspected information leakages

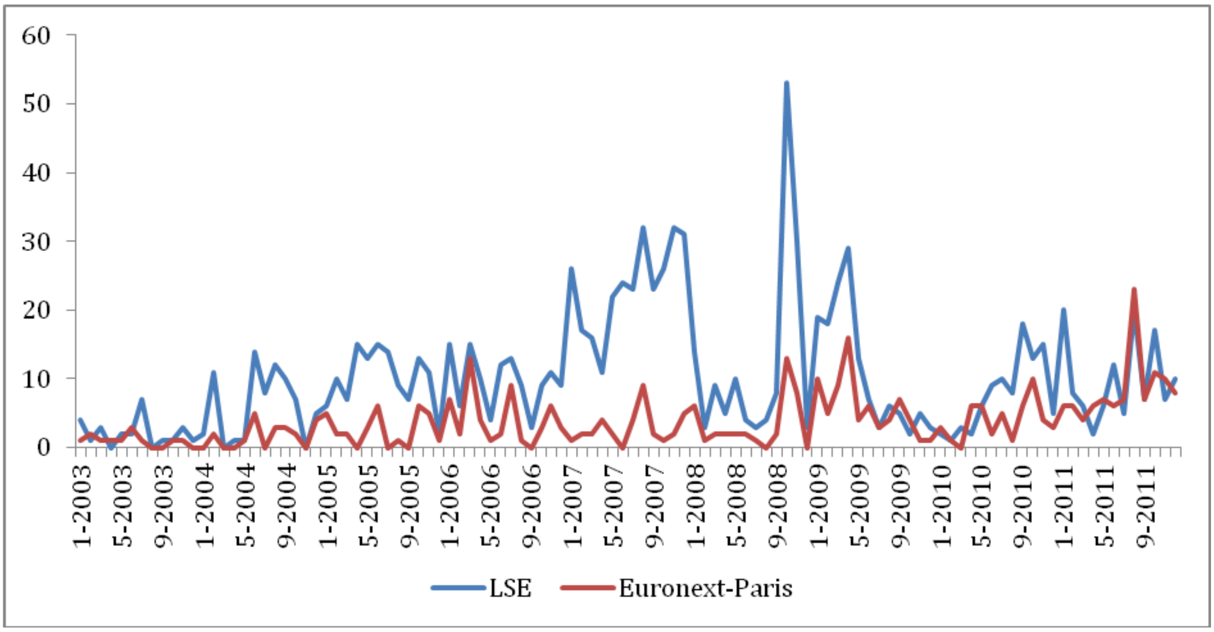

6.4.2 Market manipulation

Market manipulation involves creating a false or misleading representation of the possibility of undisclosed information with the intent to affect the market price. Market manipulation reduces the average, long-run informational efficiency of stock prices and degrades overall market quality. There are many types of market manipulation and distinctions need to be drawn between actual market manipulations and activities that may look like manipulations, such as temporary liquidity imbalances. Brogaard (2010) and Jarnecic and Snape (2010) both find that HFT participants are inclined to end the trading day in a stock-neutral position. This results in elevated HFT participation towards the end of the trading day. This, coupled with Brogaard's (2010) finding that the trading patterns of HFT's are similarly correlated could create a false impression that markets are being manipulated.

In this report, we estimate market manipulation by examining the number of suspected instances of dislocation of the end-of-day price. Dislocating the end-of-day price involves influencing closing prices so that they no longer represent the true forces of supply and demand. The determination of the closing price is fundamentally critical and is used by a range of market participants for a number of different purposes. Potential motivations for manipulating end-of-day prices include:

---

[18] Due to the availability of public company announcements we are only able to construct the information leakage metric from June 2003 for Euronext and June 2007 for the LSE.

- To modify the value of managed funds. Fund managers may dislocate the closing price in an attempt to alter the appearance of their performance and increase their ranking relative to competitors. This may be particularly associated with end-of-month or end-of-quarter reporting periods;
- To profit from derivatives positions in the underlying stock;
- To obtain a favourable price in pre-arranged off-market trades;
- To alter their customers' inference of broker execution ability;
- To maintain a stock's listing on an exchange with minimum price requirements;
- To gain inclusion in an index near stock index rebalancing days and;
- To avoid margin calls

However, not all abnormal closing prices are the result of deceitful trading strategies. Some reasons why stock prices may naturally close at unusual levels include:

- Announcements or changes in underlying instruments near the close or during the closing auction may cause large price movements;
- Brokers with a mandate to sell certain quantities of stock may be forced to become aggressive in order to liquidate before the end of the trading day;
- Some market participants may not like to hold inventory overnight and are obliged to liquidate at the close, irrespective of price;
- Participants who enter large market orders at the closing action may unwittingly cause large price movements if they are not mindful of the indicative closing price or the amount of depth available on the opposite side of the book.

Distinguishing between the various forms of abnormal end-of-day prices is a challenging task. In this report, attempts to mark the close are measured as abnormally large end-of-day price changes which exceed predetermined stock-specific thresholds. For each stock and trading day, the price change of the last 15 minutes of trading is compared to a distribution of historical price changes, which occurred during the last 30 trading days.

$$\Delta EOD_{i_t} = \frac{P_{eod,i_t} - P_{eod-15m,i_t}}{P_{eod-15m,i}}$$

$$\overline{\Delta EOD_i} = \frac{1}{30} \sum_{t=-31}^{t=-1} \Delta EOD_{i_t}$$

Where ΔEOD is the return between the closing price and the price 15 minutes prior to the close, $\overline{\Delta EOD}$ is the average return over a rolling window of thirty trading days prior to the day being analysed and σ_i is the standard deviation of ΔEOD_i over the same period. Manipulative behaviour is suspected when an end-of-day price change exceeds 3 standard deviations above or below the mean of the distribution of prior observations.

$$Potential\ Positive\ Manipulation\ if\ R_{i_t} > \overline{\Delta EOD}_i + 3 * \sigma_i$$

$$Potential\ Negative\ Manipulation\ if\ R_{i_t} < \overline{\Delta EOD}_i - 3 * \sigma_i$$

End-of-day prices that are not the result of the genuine forces of supply and demand are likely to exhibit next-day price reversion. Instances of abnormal end-of-day price changes which are followed by a price reversion of 50% or more on the open of the next trading day are considered successful attempts at marking the close (MTC). These are referred to as instances of "dislocating the end of day price" ("EOD price"). Without this reversion, the abnormal change in the end of day price could be due to changes in the true supply and demand for the stock. In the absence of this reversion, we do not consider a manipulation to have occurred.

Figure 4 shows the number of successful attempts at dislocating the EOD price for the Euronext and LSE markets respectively. Both the LSE and Euronext Paris show some variation across time, with our indicator showing elevated signs between 2005 and 2007. Since this period of heightened EOD price manipulation, the LSE has reverted to its historically low levels, whilst Euronext Paris displays some infrequent periods of price dislocation.

Figure 4. Number of successful EOD manipulations

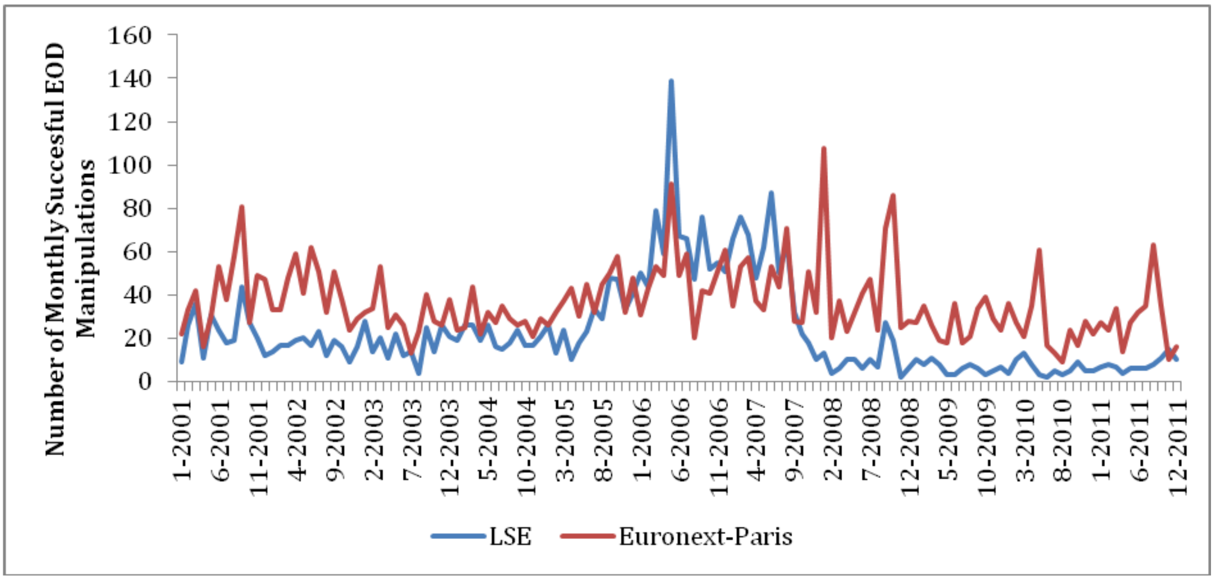

7. Research design

7.1. A simultaneous structural equations model for market quality research

To study the cross-sectional determinants of HFT participation, it is necessary to include variables that are publically available over a long time series. The modelling considers the fact that quoted spreads and measures of market integrity are simultaneously determined. Higher suspected cases of market manipulation raises volatility and reduces order aggressiveness, leading to higher bid-ask spreads. Conversely, quoted spreads are a non-trivial execution cost of market manipulation. Higher quoted spreads, therefore, reduce the incidence of manipulation, ceteris paribus. Market integrity, efficiency, and the level of high-frequency trading are similarly, simultaneously determined. The lower cost of trading increases the level of high frequency trading, as does improved market integrity. On the other side, an increase in the level of high-frequency trading potentially leads to lower bid-ask spreads if high-frequency trading leads to a more competitive market environment characterised by improved market integrity. We adopt a three-stage least squares (3SLS) method in the belief that efficiency (spreads), integrity (EOD manipulation or Information leakage) and HFT participation residual errors will likely be cross-equation correlated.

The empirical model structure is a simultaneous set of three structural equations describing market integrity (AI – the number of alert instances), market efficiency (SPR – the effective spread efficiency), and a specific market design change (HFT – HFT participation):

Integrity Eqn:

AI_i = ($\widehat{Spr}, \widehat{HFT},$ Control variables, Market Design Changes, Fixed effects) (i)

Efficiency Eqn:

SPR_i = ($\widehat{AI}$, $\widehat{\mathbf{HFT}},$ Control variables, Market Design Changes, Fixed effects) (ii)

Market Design Eqn:

HFT_i = ($\widehat{AI}$, $\widehat{Spr}$, Control variables, Market Design Changes) (iii)

Where;

$AI_{i,t}$ = A constructed variable including the number of suspected cases of insider trading ("Insider") and market manipulation ("MTC")

$Spr_{i,t}$ = A constructed variable of effective spreads. This is a measure of the cost of a round trip transaction.

$HFT_{i,t}$ = A variable that is constructed to proxy the level of HFT participation in the market using the order to trade and cancel to trade ratios[19].

The instrumental variables for the endogenous variables are based on reduced-form equations of all the exogenous and predetermined regressors; hat symbols signify these predicted values. Because each of the endogenous variables could in principle affect each other, we assure the order condition for identification by excluding from each equation two control variables present elsewhere in the system. In each equation, the excluded variables are control variables found to be insignificant in preliminary single-equation estimations of the focal equation but highly significant in the other two structural equations.

We hypothesise the following regression relation between the level of integrity of an exchange (proxied by either the percentage of suspected cases of insider trading or the level of marking around the close) and a range of explanatory variables described below:

$$AI_t = \alpha + \beta_1 \widehat{SPR}_t + \beta_2 \widehat{HFT}_t + \beta_3 Std_t + \beta_4 Ret_t + \beta_5 Vol_t + \beta_6 AI_{t-1} + \beta_7 MiFID + \beta_8 Quarter + \varepsilon' \quad \text{(i')}$$

For the level of market efficiency (proxied by the size of the quoted spread):

$$SPR_t = \alpha' + \beta_9 \widehat{AI}_t + \beta_{10} \widehat{HFT}_t + \beta_{11} Price_t + \beta_{12} Vol_t + \beta_{13} Std_t + \beta_{14} SPR_{t-1} + Yearly + \varepsilon'' \quad \text{(ii')}$$

And for our HFT participation equation:

$$HFT_t = \alpha'' + \beta_{15} \widehat{AI}_t + \beta_{16} \widehat{SPR}_t + \beta_{17} Std_t + \beta_{18} Vol_t + \beta_{19} HFT_{t-1} + \beta_{20} HFT_{t-2} + \varepsilon''' \quad \text{(iii')}$$

Where;

Std_t = Mean standard deviation of daily returns at time period *t*.

Ret_t = Mean daily return at time period *t*.

Vol_t = Mean turnover at time period *t*.

$MiFID$ = Dummy variable for the introduction of MiFID (Equals 0 prior to November, 2007, and 1 subsequent).

$Quarter$ = Dummy variable to account for the end of each financial quarter (=1 for March, June, September and December, and 0 for all other months).

$Price_t$ = Mean price at time period *t*.

[19] The order condition for identification is satisfied in the following ways: Integrity is identified through the exclusion of the Yearly Fixed Effects and Price variables; Efficiency is identified through the exclusion of the End of Quarter Effects and market return, whilst HFT is identified by excluding the End of Quarter Effects, the Yearly Fixed Effects and total Price.

Yearly = Yearly fixed effects are applied to the constructed regressions.

ε', ε'', ε''' = residual error terms.

The specifications are modelled separately for all LSE and Euronext Paris securities. Beyond accounting for the potential simultaneity bias in the multiple-equation results, we address the potential cross-equation correlation of the error terms by estimating using three-stage least squares. In the regression output we report the coefficients and standard errors for each of the explanatory variables. ***, **,* indicate the level of significance at 1%, 5%, and 10%, respectively.

7.2. Stationarity of our estimates and adjustment for Autocorrelation

Durbin-Watson tests indicated the presence of autocorrelation. In order to assess the optimal lag structure we used general to specific F-tests, starting with six lags. The optimal lag structure was found to be 1 month for our Spread, Insider Trading and Marking the Close metrics. Our HFT proxy was found to exhibit a two month lag structure. Our regression analyses have used these optimal lag structures to minimise the autocorrelation, with reported Durbin-Watson statistics for each of our regressions insignificant at the 5% level. The parameters themselves are adjusted for possible auto-correlated disturbances by adopting Newey-West standard errors.

8. Results

8.1. Descriptive statistics

For each of the metrics used as dependent variables in the equations (i), (ii) and (iii) descriptive data is presented in Table 1. These descriptive statistics are split both by market and into time periods that are pre- and post- the introduction of MiFID. This distinction has been made due to the significant structural changes MiFID initiated, including the fragmentation of the European equities market.

Table 1. Dependent variable descriptive statistics by market

	Minimum	Mean	Median	Maximum	Std. Deviation
Panel A - LSE					
Marking the Close (number)					
- Pre-MiFID	4.00	31.85	24.00	139.00	22.95
- Post-MiFID	2.00	7.55	7.00	27.00	4.62
Information Leakage (%)					
- Pre-MiFID	0.00	24.47	25.00	60.00	31.56
- Post-MiFID	3.77	20.74	17.86	60.00	13.52

	Minimum	Mean	Median	Maximum	Std. Deviation
Effective Spread (basis points)					
- Pre-MiFID	10.31	17.43	14.84	33.31	6.07
- Post-MiFID	8.13	10.82	10.69	16.45	2.11
Cancel to Trade					
- Pre-MiFID	0.52	3.07	3.20	6.53	1.39
- Post-MiFID	2.20	17.91	15.68	39.50	11.13
Panel B - Euronext Paris					
Marking the Close (number)					
- Pre-MiFID	13.00	39.34	37.00	91.00	14.34
- Post-MiFID	9.00	31.81	27.00	108.00	18.72
Information Leakage (%)					
- Pre-MiFID	0.00	25.35	24.50	100.00	23.54
- Post-MiFID	0.00	28.84	26.09	100.00	18.04
Effective Spread (basis points)					
- Pre-MiFID	5.91	11.71	11.26	21.33	3.12
- Post-MiFID	5.28	7.68	7.53	12.23	1.71
Cancel to Trade					
- Pre-MiFID	0.91	1.42	1.31	2.79	0.32
- Post-MiFID	1.70	8.01	6.17	25.93	6.28

The introduction of MiFID appears to have reduced the prevalence of marking the close violations on the LSE significantly, with smaller reductions also observed on Euronext. The impact of MiFID on the amount of information leakage is less clear, with minor reductions observed on the LSE but minor increases observed on Euronext. Our proxy for market efficiency, effective spreads, exhibits significant reductions in the post-MiFID period. This is likely driven by the increased fragmentation observed in the market subsequent to the introduction of MiFID. Our proxy for HFT, the cancel to trade ratio, exhibits a marked increase subsequent to the introduction of MiFID, increasing by over 5 times from its pre-MiFID levels for both the LSE and Euronext. The time series variation in each of the metrics documented in Table 1 is displayed and discussed in the remainder of this section.

One of the hallmarks of HFT participants is the ability to update quotes on an almost instantaneous basis. This updating is necessary to minimise the adverse selection risk faced by HFT participants acting as market makers, and also to take advantage of opportunities created by the lower levels of latency faced by HFT participants. The message to trade ratio is used by Hendershott, Jones and Menkveld (2011) as a proxy for the level of HFT participation in the market. Figure 5 displays the Level 1 order to trade ratio for the LSE and Euronext Paris, a publically available measure of the message to trade ratio.

The order to trade ratio at the first level is calculated as the number of updates to the best bid and ask divided by the total number of trades. This ratio remains within a very narrow band of between 2 and 5 up until the end of 2007. The implementation of MiFID in November 2007, in addition to the introduction of co-location and direct market access around this time removed significant structural barriers for a range of market participants, allowing better access to a range of markets. The practice of HFT which normally involves significant quote posting and revisions in part explains the steady increase of this ratio up until the beginning of 2010. The order to trade ratio reaches its high point for both exchanges during August 2010, reaching approximately 15 and 25 for the LSE and Euronext Paris, respectively. Particularly noteworthy is that the dramatic increase in the order to trade ratio through to 2010 is followed by a more rapid decline after this time. This coupled with the diverging ratios after 2010 provide interesting results that warrant further investigation.

A very strong correlation exists between the top of book order to trade ratio documented in Figure 5 and the 10 level order to trade ratio documented in Figure 6. The 10-level order to trade ratio is constructed in a similar manner to the top of book ratio, with the key difference being that all changes to the limit order book within the top 10 levels of depth are taken into account. This metric is sampled from 2001 and similar to Figure 5 shows little variation in the first seven years. This ratio doubles during 2007 and continues this momentum through to 2008. The spike in order to trade ratios for the LSE and Euronext Paris observed in Figure 5 in 2010, is similarly reflected in the Level 10 data with a high of 200 cancels for every trade being reached in August 2010. Overall, the veritable explosion in quote traffic is indicative of a significant amount of high frequency trading occurring in Europe.

Figure 5. Euronext Level 1 order to trade ratio 5 day moving-average

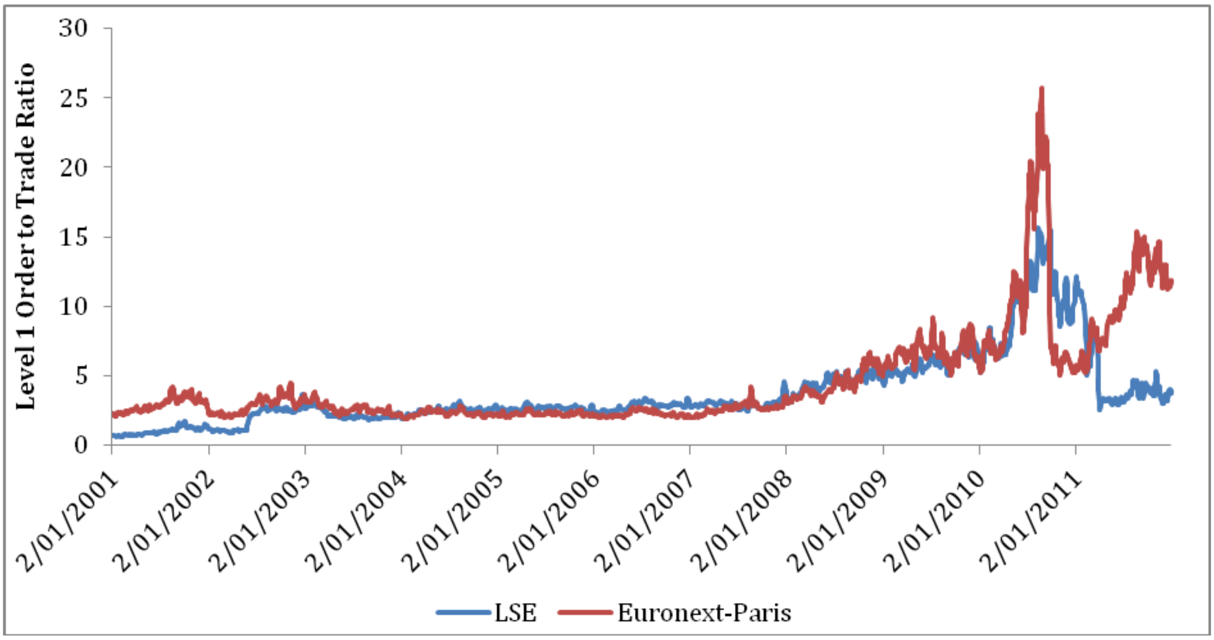

Figure 6. Level 10 order to trade ratio 5 day moving-average

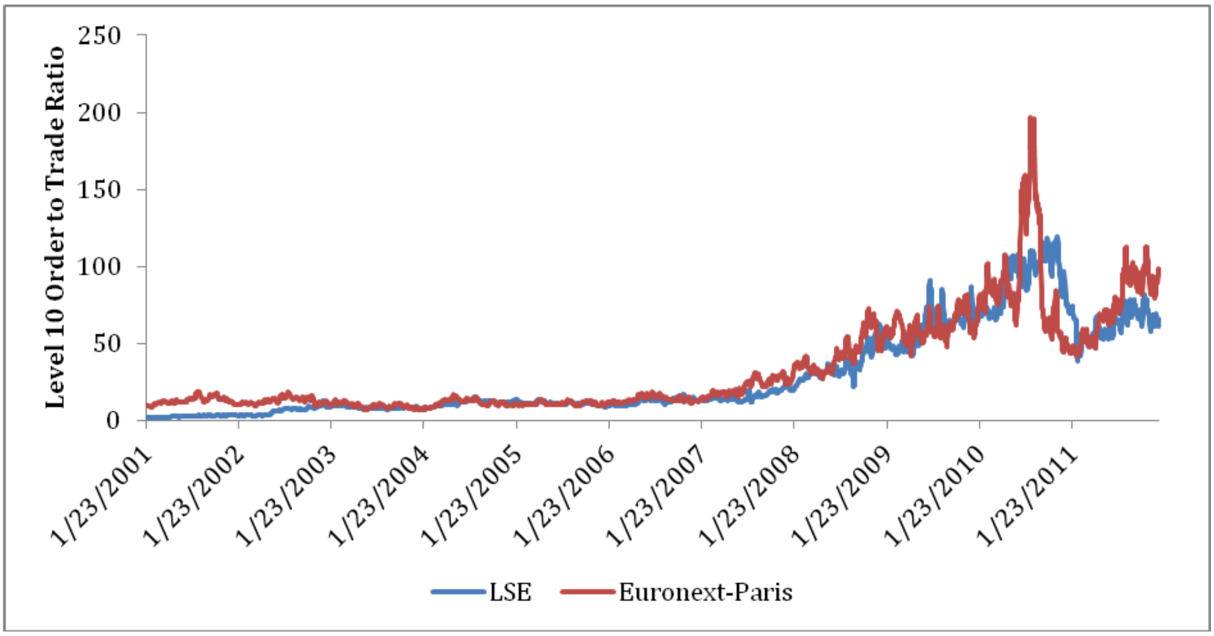

Figure 7 shows the proportion of cancelled orders to trades for the LSE and Euronext Paris. If number of orders increase relative to the number of executed trades, as is suggested in the previous set of figures, then it is likely that the number of cancelled trades will also increase. Figure 7 shows little deviation in the cancel to trade ratio between 2001 and 2007. As the proportion of orders increases, this ratio also increases significantly for both the LSE and Euronext Paris, reaching a high point in 2010. There is a reverse in the trend after this time with the cancel to trade ratio falling significantly throughout 2011. This pattern is observed for both exchanges in Europe. The causes of the sharp decline warrant further investigation. This ratio will be used as our proxy for HFT activity, in accordance with Menkveld (2012).

Figure 7. Cancel to trade ratio 5 day moving-average

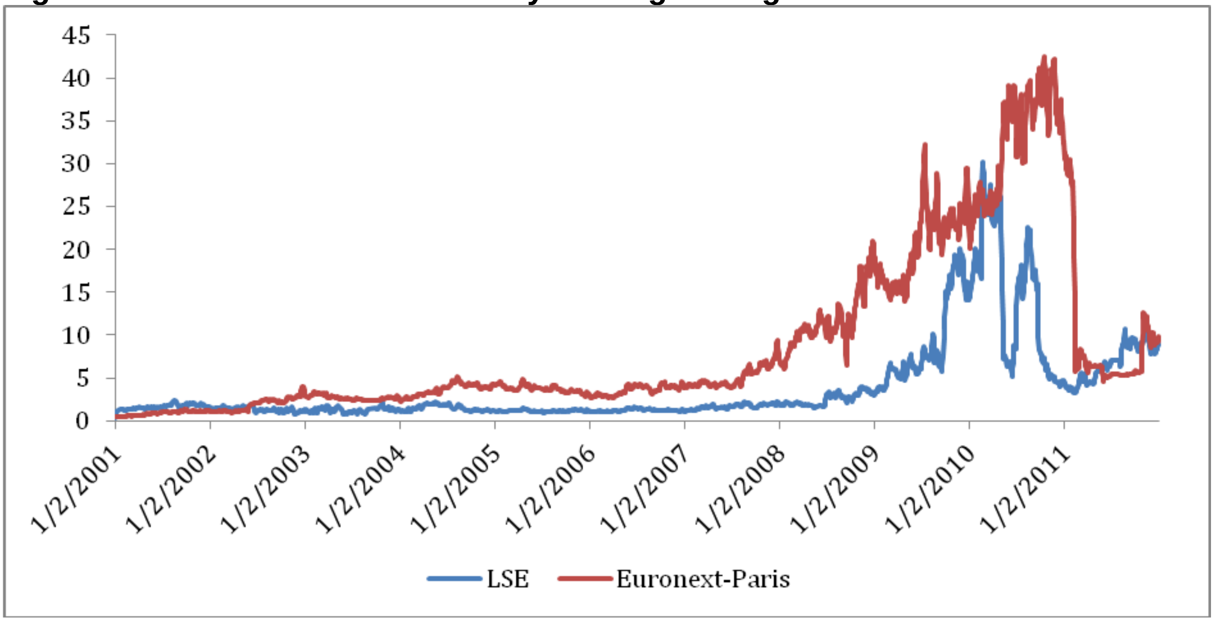

Cvitanic and Kirilenko (2010) predict that the introduction of HFT participants will reduce both the volume and value of individual trades, due to their desire to minimise both their inventory holding and adverse selection costs. Figure 8 and Figure 9 document the evolution of these variables on the LSE and Euronext over the period January 2001 to January 2012.

Figure 8 shows the average trade size (shares) of LSE and Euronext Paris securities between 2001 and 2011. The figure indicates a dissimilar trend up until 2005, where the LSE experienced a decline in the average trade size, whilst Euronext Paris followed an upward trend, with average trade size more than quadrupling over this period. Following this time, the figure shows a steady downward trend for both the LSE and Euronext Paris beginning in 2005 before a plateau is reached in 2008. The average trade size is approximately 5000 and 400 shares for the LSE and Euronext Paris, respectively. This pattern substantiates the predictions of Cvitanic and Kirilenko (2010).

Figure 8. Average trade size (shares) 5-day moving-average

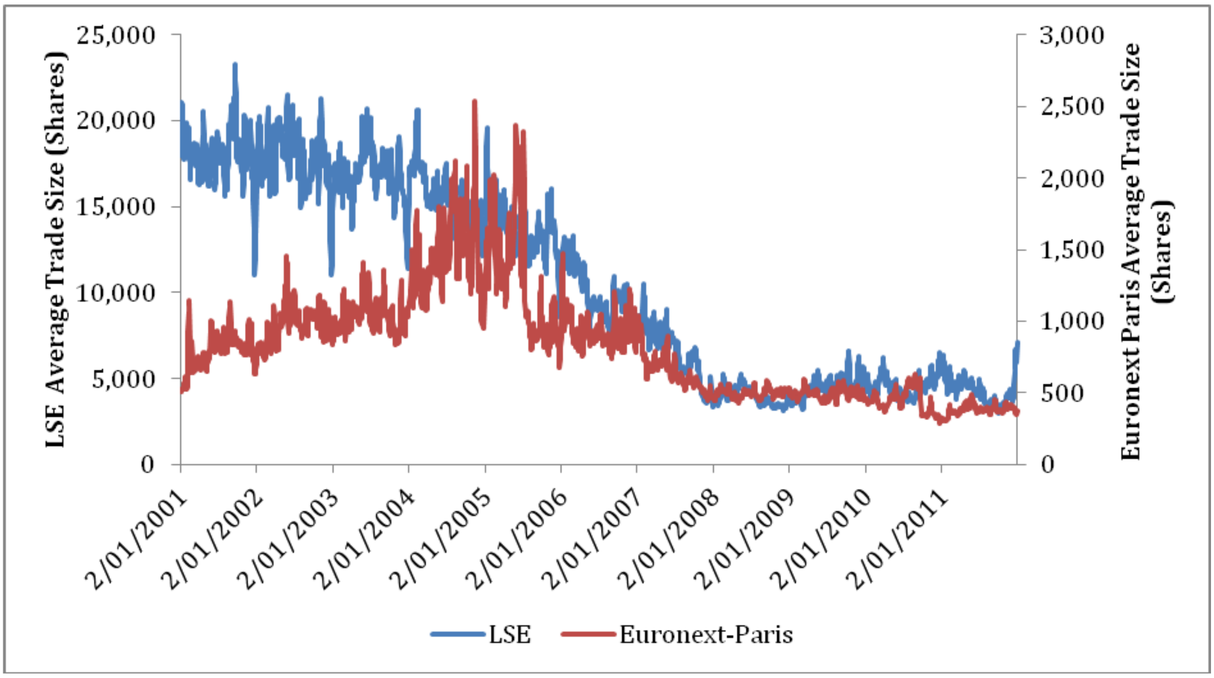

Figure 9 shows average trade size (value) for the LSE and Euronext-Paris between 2001 and 2011. The figure illustrates a similar movement in average trade values for these exchanges across time. As with traded volume, a systemic change occurs in 2007, with average trade size declining from around £30,000 (€20,000) in January 2007 to approximately £9,000(€9,000) at the beginning of 2009 for the LSE (Euronext-Paris). From 2009 to 2011 both the LSE and Euronext Paris exhibit a slower rate of change, with average trade sizes falling further to approximately £5000 and €6000, respectively. Whilst it is possible that the reductions in volume are driven by a systematic increase in share prices over the sample period, it does not account for the concurrent reductions in traded value nor the structural breaks observed in 2007.

Figure 9. Average trade value 5 day moving-average

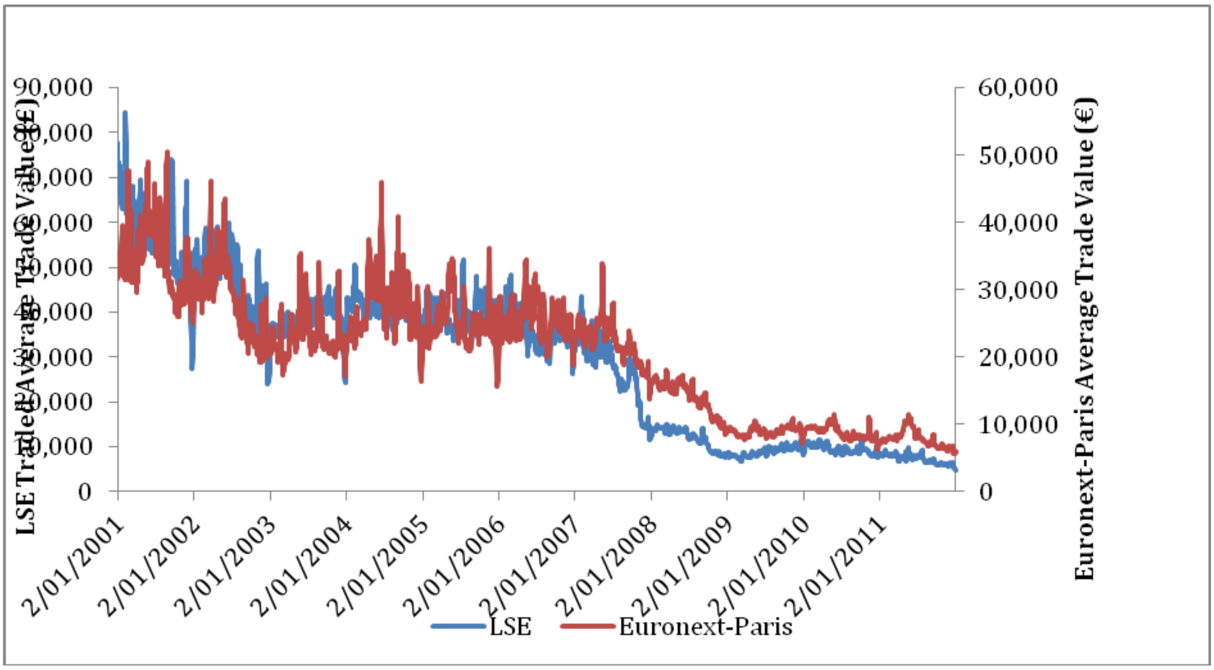

Figure 10 shows effective bid-ask spreads for the LSE and Euronext Paris from 2001 to 2011. Both Brogaard (2010) and Jarnecic and Snape (2010) find that HFT participation results in a reduction in the bid-ask spread. This outcome is similarly modelled by Jovanovic and Menkveld (2010). The fragmentation of both European and global equities markets combined with technological innovation has seen a significant reduction in order processing costs that are normally recovered through the bid-ask spread. In addition, the empirical work of Hendershott, Jones and Menkveld (2011) suggest an increase in HFT participation is associated with a reduction in bid-ask spreads because HFT traders are better able to manage information asymmetries. Figure 10 shows that the effective spread has fallen through time. Relative to 2001, effective spreads have fallen by approximately 60% and 45% for the LSE and Euronext Paris, respectively. Given the constraint of a minimum tick size, there would appear to be little room for bid-ask spreads to fall further as is evidenced towards the end of the sample period.

Figure 10. Euronext monthly effective spreads

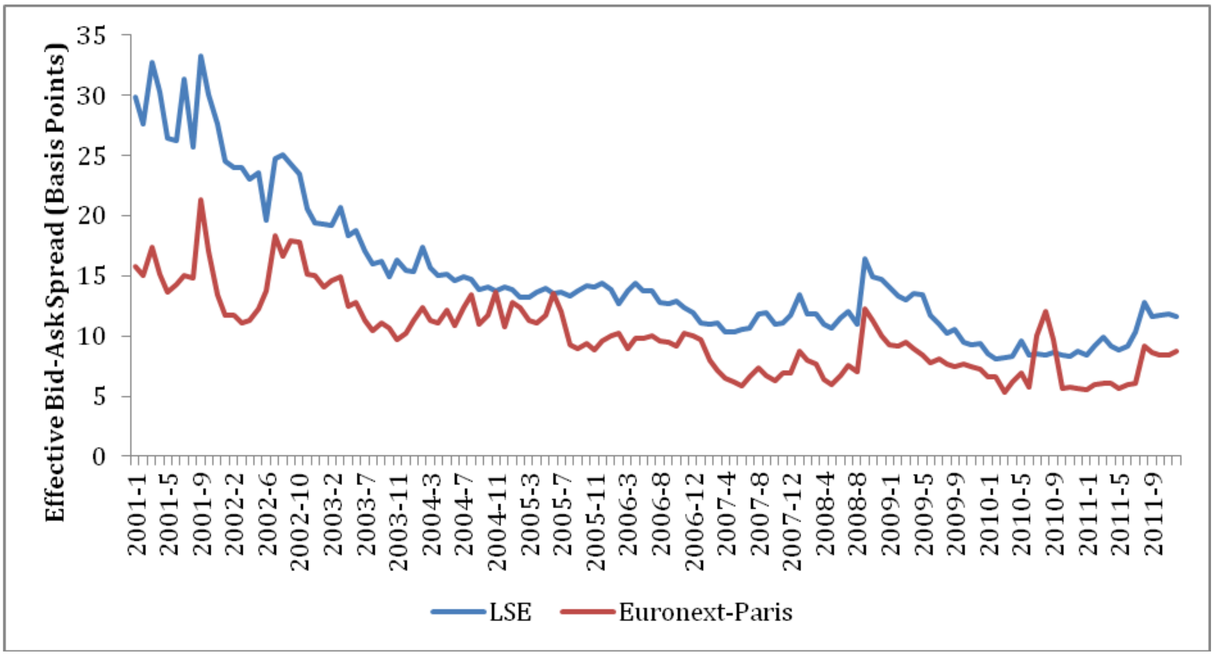

In their analysis of NASDAQ trade and quote data, Hasbrouck and Saar (2010) identify periodicities in their 2008 and 2009 data. These periodicities exist in "wall-clock" time - that is taking timestamps which define milliseconds since midnight. This deconstructs each second into 1000 equally sized buckets, and analyses the arrival rate of orders based upon the part of the second in which they were entered. If there were no patterns to order entry, approximately $1/1000^{th}$ of orders would exist within every millisecond. Instead, Hasbrouck and Saar (2010) identified very frequent order entries immediately following the start of the second, and less strong though prevalent order entries immediately following the half-second barrier.

They attribute these peaks in order entry activity to automated trading systems that periodically access the market, near the second and half-second boundaries. They find that these periodicities are on much longer intervals than the minimum latency within which HFT participants are able to react to orders, which is in the sub 20ms range. This leads Hasbrouck and Saar (2010) to identify two subsets of HFT participants, those that programmed to access and revisit the market periodically, potentially in order to "slice and dice" an institutional order, and those that react to information events. The works of Roll (1985) and others show that where liquidity is demanded, both informed and discretionary uninformed traders will congregate. It appears that these longstanding theories are also applicable to the millisecond environment.

Figures 11 and 12 document the change in millisecond level order entry at two separate points in time. Figure 11 considers the month of January 2007, being immediately prior to the introduction of MIFID, which facilitated the market fragmentation thought to have accelerated HFT participation. An equal incidence of order arrivals through time would result in a uniform distribution at 0.1% (that is, 1/1000 * 100%). This is very closely represented by the figures for each market in 2007. Whilst the LSE exhibits some bulge in order entries between 300 and 550ms these are not sharply defined.

Figure 11. Euronext millisecond remainders January 2007

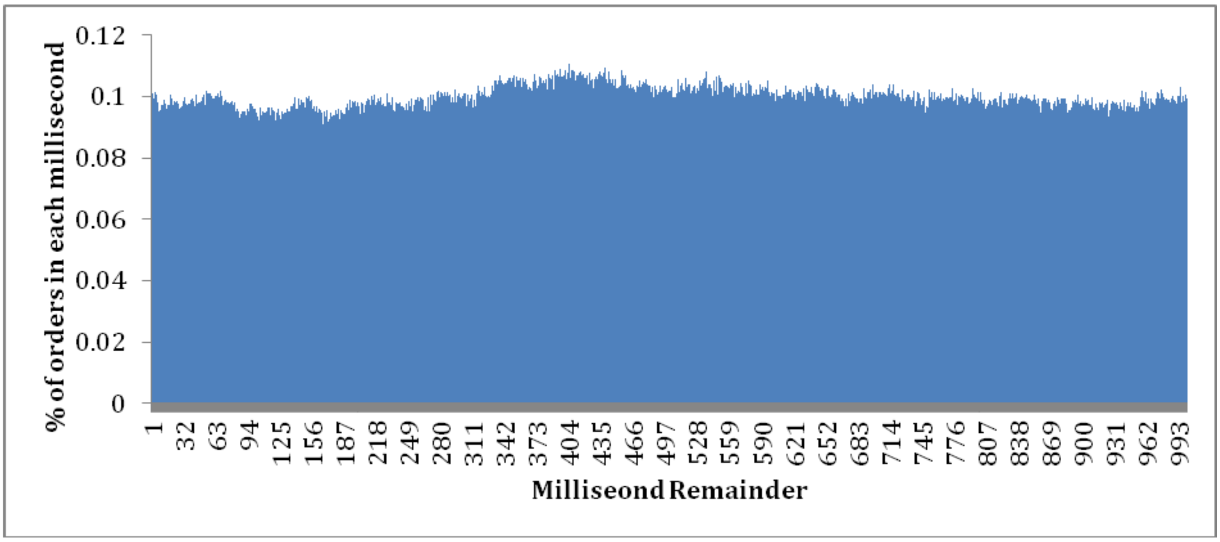

Figure 12. LSE millisecond remainders January 2007

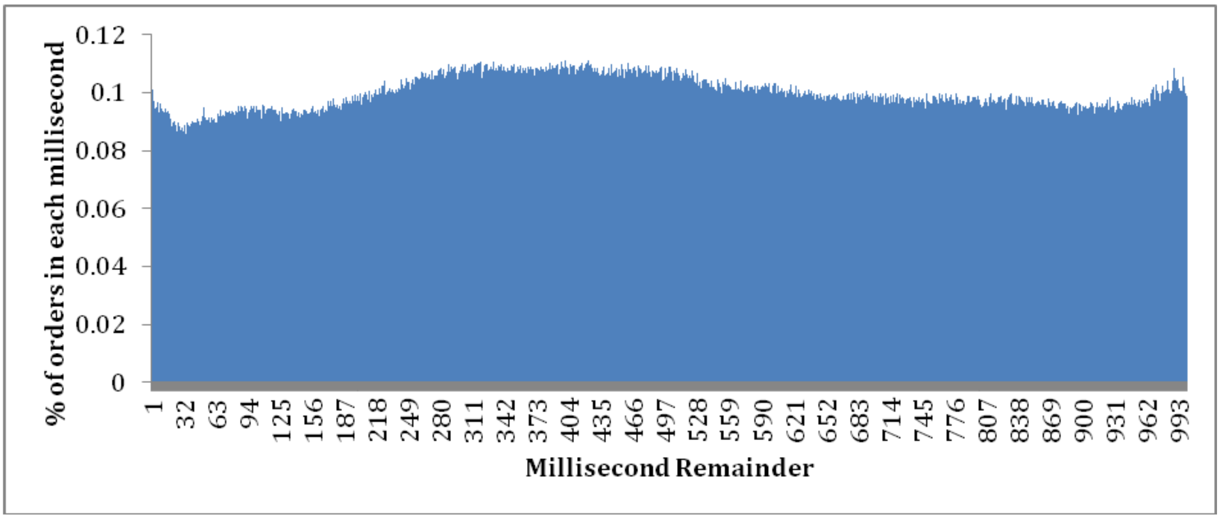

Figure 13 shows the millisecond order and trade entry pattern for Euronext in February 2012. Euronext exhibits very marked periodicities around 150ms and 300ms. The shape of these periodicities is similar in nature to the double peak observed by Hasbrouck and Saar (2010) in their NASDAQ data for the periods 2008 and 2009. Hasbrouck and Saar (2010) attribute this double peak to clustering either in transmission time (due to geographic latency, potentially due to co-location) or potentially to other technological challenges associated with handling large volumes at the firm level.

Figure 13. Euronext millisecond remainders February 2012

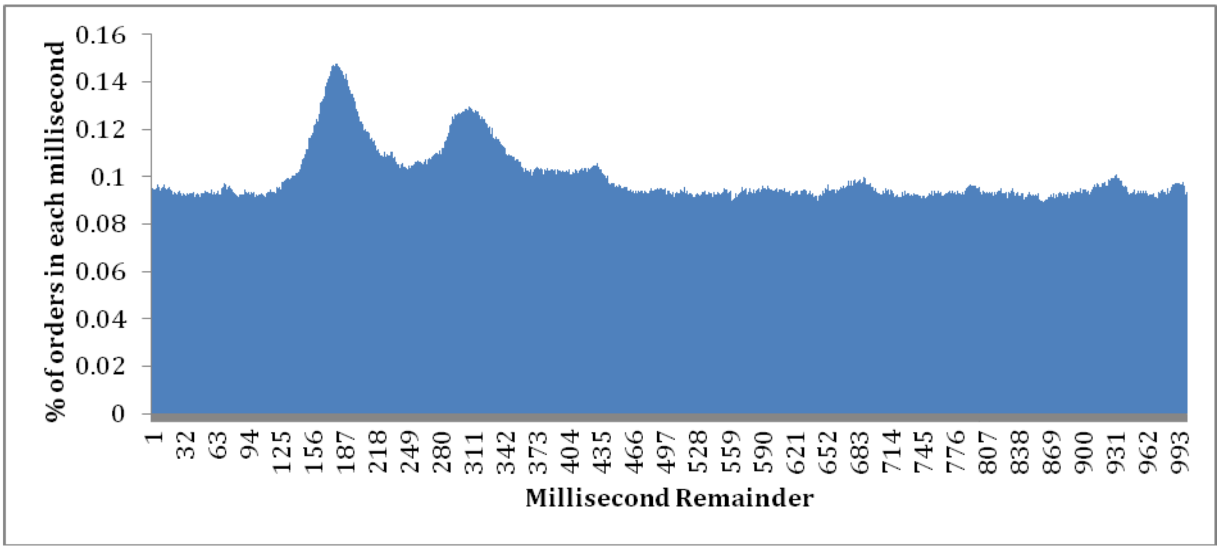

Figure 14 for the LSE from February 2012 indicates an even more striking periodicity. There are ten double peaks, offset from the second boundary by approximately 20ms, with each of the second peaks approximately 50ms away from the first. These ten double peaks appear to occur just after the 100ms boundary. This appears to be evidence of the same kind of algorithmic approach to trading found on the NASDAQ and LSE, but occurring every 100ms. As in the Hasbrouck and Saar (2010) 10-second analysis, there appears to be a heightened order submission rate around the second, first and third 100ms periods.

Figure 14. LSE millisecond remainders February 2012

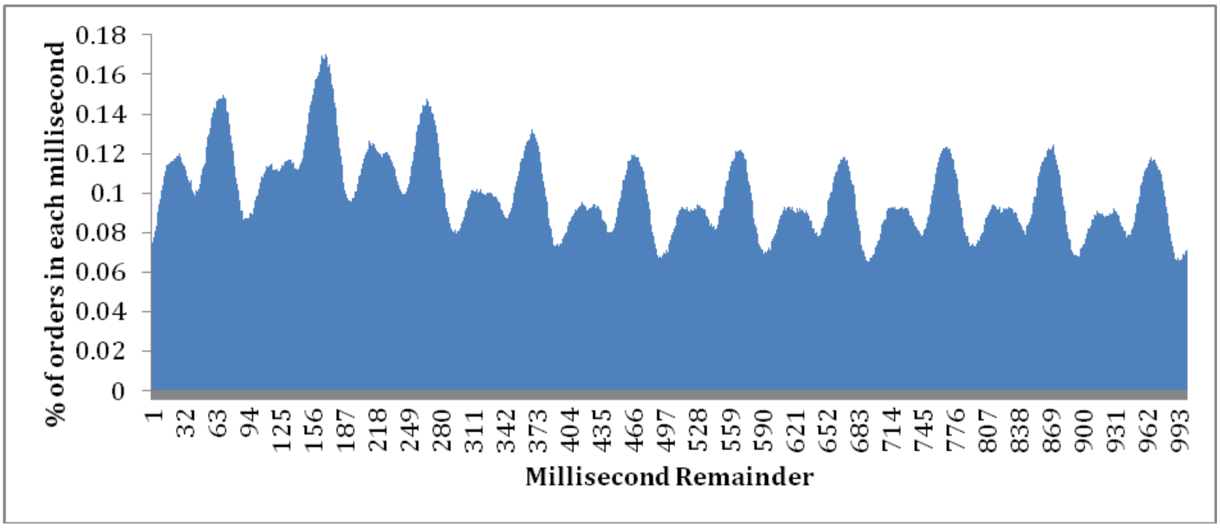

The contrasting appearance of millisecond remainders between 2007 and 2012 indicates that there now exist high frequency trading patterns that did not exist at the beginning of 2007. This evidence is consistent with the empirical findings of Brogaard (2010) and Jarnecic and Snape (2010) that find significant levels of HFT participation in the LSE market using 2009 and 2010 data.

8.2. Regression results

Table 2 presents the results of the 3SLS estimation described in the research design. Specifically, the determinants of effective quoted bid-ask spreads are examined for the LSE and Euronext Paris. [20] We find that the percentage of potential insider trading violations (our proxy for market integrity) is positively and statistically significant across both the LSE and Euronext markets, which suggests that a deterioration of market integrity is associated with higher trading costs. This result is consistent with the premise that market efficiency can be improved with appropriate controls and supervision of the trading behaviour of market participants. Whilst our first measure of market integrity is significant for both Euronext and the LSE, our second measure, the dislocation of end-of-day price is only weakly associated with bid-ask spreads on the LSE. Further, it does not have a significant impact on bid-ask spreads on Euronext. This result is consistent with Comerton-Forde and Putnins (2011) who show the bid-ask spread is likely to be lower on days where end-of-day manipulation is present, but higher the following day as market participants react to the possibility of market manipulation.

Table 2. Determinants of the effective spread

	Euronext		LSE	
Intercept	-0.217	-0.118	1.269***	0.523**
	(0.147)	(0.159)	(0.278)	(0.200)
$\widehat{Number\ of\ Successful\ MTC}$	0.000	-	0.001*	-
	(0.000)	-	(0.000)	-
$\widehat{\%\ Information\ Leakage}$	-	0.047**	-	0.117***
	-	(0.019)	-	(0.011)
Ln(Price)	-0.034***	-0.042***	-0.119***	-0.082***
	(0.007)	(0.008)	(0.021)	(0.011)
Ln(Traded Volume)	0.019**	0.018**	-0.010	0.001

[20] Note that due to the restricted availability of announcement data from January 2003 our specifications have been run with only 2003-2011 data. Unreported specifications using the MTC proxy were run over the period 2001 – 2011 with similar results. These results are available upon request from the authors.

	Euronext		LSE	
	(0.007)	(0.007)	(0.009)	(0.008)
Ln(Volatility)	0.003	0.008**	0.012***	-0.001
	(0.003)	(0.003)	(0.003)	(0.002)
$\widehat{HFT}$	-0.009**	-0.014***	-0.045***	-0.012**
	(0.004)	(0.004)	(0.006)	(0.003)
MiFID	-0.004	-0.005	0.0318**	0.001
	(0.005)	(0.010)	(0.013)	(0.007)
$Lag_1\left(\widehat{Spread}\right)$	0.632***	0.470***	-0.004	0.288***
	(0.091)	(0.072)	(0.085)	(0.073)
N	105	105	105	105
Durbin-Watson	1.69	1.76	1.84	1.62
Adjusted R-Squared	0.805	0.787	0.864	0.86
Year Fixed Effects	Yes	Yes	Yes	Yes

Results from our regression analysis also suggest that HFT participation has a negative and significant impact for bid-ask spreads on LSE and Euronext Paris listed securities. This result is foreshadowed by the intense competition for order-flow and market-making revenues documented in the literature. The introduction of MiFID, which has resulted in greater competition for incumbent exchanges, largely generated by an increase in HFT activity, has resulted in lower trading costs for market participants.

The opposing signs on our coefficient for MiFID, whilst initially a curious result, may be explained by differences in market structures prior to 2007. The introduction of the passporting rule by MiFID allowed stocks listed in any EU country to be traded in any of the venues also regulated by an EU country. It thus provided for significant competition in the trading of securities which is likely to have had a larger impact on Euronext-Paris than the LSE market because the UK market was already subject to competition by this time.

Our control variables for price behave as predicted, with higher priced securities reducing the spread as recorded in basis points across all specifications. The results for traded volume are positive and significant for the Euronext sample. This result is likely driven by the market-wide nature of the estimation, with larger traded volume during points of market stress, such as the GFC. Where volatility is found to be significant it increases spreads, as anticipated.

Included year fixed effects are found to be significant and negative, reflecting the deterministic downward trend in spreads through our estimation period attributable to MTFs like Instinet-Europe and to MiFID-triggered competition for order flow across exchanges.

Finally, the lag of the bid-ask spread is found to be positively significant in three of the four specifications. This result is driven by the presence of first-order autocorrelation. The Durbin-Watson statistics ranging from 1.62 – 1.84 are not found to be statistically different from 2 at the 5% level, indicating that the inclusion of the first order lag has accounted for the detected autocorrelation of our residuals.

Table 3 presents the results for market integrity of the 3SLS estimation described in the research design section. Specifically, the determinants of market integrity are examined for the LSE and Euronext Paris. Our results firstly indicate that HFT participation is negatively and significantly related to market integrity, which suggests that HFT has not only created a more competitive market environment resulting in lower trading costs, but also improved the integrity of the operating environment.

The coefficient on the cost of transacting variable is positive and significant to the number of integrity violations in three of the four specifications. This indicates a significant association between attempted market manipulations and bid-ask spreads. This relationship is well accepted by the finance literature as a result of traders and/or market makers identifying the existence of asymmetric information and becoming less willing to trade.

Table 3. Determinants of market integrity

	Euronext MTC	Euronext Insider	LSE MTC	LSE Insider
Intercept	-1.942	-0.438	-596.282***	-490.450***
	(2.248)	(2.454)	(192.204)	(136.297)
Ln(Volatility)	0.005	-0.082	0.962	10.120***
	(0.035)	(0.057)	(3.223)	(3.119)
Ln(Return)	-0.118	-0.609	-27.244	-131.842***
	(0.216)	(0.424)	(27.548)	(26.118)
Ln(Traded Volume)	0.116	0.051	24.585***	23.688***
	(0.091)	(0.108)	(7.949)	(6.055)
$\widehat{HFT}$	-0.011	-0.097***	-1.002	-5.774***
	(0.020)	(0.032)	(2.368)	(2.092)
Ln($\widehat{Spread}$)	0.386***	0.307***	-0.958	14.221**

	Euronext MTC	Euronext Insider	LSE MTC	LSE Insider
	(0.092)	(0.094)	(8.412)	(5.495)
MiFID	0.096*	0.323***	-8.021	-8.864*
	(0.056)	(0.078)	(5.555)	(4.478)
End of Quarter Fixed Effect	-0.000	0.077***	1.946	3.536*
	(0.016)	(0.023)	(2.229)	(1.891)
$Lag_1\left(\widehat{\overline{Marking\ the\ Close}}\right)$	0.001***	-	0.568***	-
	(0.000)	-	(0.071)	-
$Lag_1\left(\widehat{\overline{\%\ Insider\ Trading}}\right)$		0.003***		-0.032
		(0.001)		(0.061)
N	105	105	105	105
Durbin-Watson	1.73	1.89	2.24	1.87
Adjusted R-Squared	0.252	0.265	0.667	0.579
Year Fixed Effects	No	No	No	No

As in our other specifications, MiFID appears to have had differing impacts on the LSE and Euronext. These differences are likely to do with the level of enforcement and regulation existing in each market. In discussions as recent as March 2012, the Economic and Monetary Affairs Committee rapporteur Arlene McCarthy argued that the fragmentation which MiFID has fostered makes it more difficult to detect instances of manipulation or misconduct.[21] It is potentially this increased fragmentation that has led to an increase in the difficulty of enforcement in Paris, whilst the level of existing fragmentation in the UK may not have led MiFID to have as pronounced an impact. The fragmentation of markets leads to a greater ability of manipulators to "hide" their conduct amongst the "crowd" of multiple markets. This difficulty of cross-market manipulation may be one of the unintended consequences of the fragmentation caused by MiFID, and bears further scrutiny in its own right.

An additional potential explanation is provided by Cumming et al. (2011). They analyse the impact of the introduction of MiFID on US cross-listed ADR's and find that MiFID actually leads to a reduction in the European trading of these securities. They attribute this to the complexity of the newly introduced regulations. If the perceived complexity of the MiFiD reforms resulted in

[21] European Parliament News Release – 12th of April, 2012.

more credible enforcement at a national level in the UK than in Paris, the fungibility of instruments between markets may have led insiders and/or manipulators to prefer trading in Paris as compared to the UK.

Our end of quarter fixed effects is found to have a positive relationship with the level of integrity violations. This evidence explains in part the initial non-stationarity of our integrity results prior to inclusion of this deterministic trend variable. It is entirely consistent with incentives for individuals to commit integrity violations being greater at the end of the quarter due to the calculation of listed fund returns. Similar to our trading cost specifications, our other control variables behave in the expected fashion. Finally, the lag of the MTC and Insider variables is found to be positively significant in three of the four specifications. This result is driven by the presence of first-order autocorrelation. The Durbin-Watson statistics ranging from 1.87-2.24 now after inclusion of these lagged terms are not found to be statistically different from 2 at the 5% level.

Table 4 presents the HFT results of the 3SLS specification described in the research design section. Specifically, the determinants of HFT participation are examined for the LSE and Euronext Paris. Our proxies for market manipulation (marking the close and information leakage) are negatively associated with the degree of market participation by high-frequency traders. We interpret this result to mean that high-frequency traders will avoid trading in venues that are characterised by lower levels of market integrity or more likely that they inhibit the practice by their presence in the market, particularly at the close. Our trading cost instrument further shows that as the cost of trading increases, there is a resultant negative impact on the level of HFT activity. As described in a previous section on the drivers of HFT, high-frequency traders are particularly sensitive to rising trading costs. HFT participation does not appear to be significantly related to the total amount of volume in the market. This lack of relationship, despite empirical evidence of HFT participants preferring larger stocks, is likely due to the market-level aggregation conducted in our study.

Table 4. Determinants of HFT participation

	Euronext		LSE	
Intercept	-115.016***	-27.037	39.762	126.759
	(40.714)	(41.695)	(33.882)	(129.012)
Ln(Volatility)	2.452***	-0.898	0.872*	4.845**
	(0.798)	(0.928)	(0.522)	(1.998)
Ln(Return)	-33.674***	-9.170	0.34	18.999
	(8.761)	(7.716)	(4.817)	(18.647)
$\overline{Number\ of\ Successful\ MTC}$	-0.258***	-	-0.013	-
	(0.048)	-	(0.017)	-

	Euronext		LSE	
$\widehat{\%\ Information\ Leakage}$	-	-20.523***	-	-27.459***
	-	(3.579)	-	(9.620)
Ln(Spread)	-4.611***	4.037**	-2.641**	-18.553***
	(1.469)	(1.601)	(1.294)	(5.043)
Ln(Volume)	5.473***	1.776	-1.657	-5.531
	(1.855)	(1.857)	(1.403)	(5.259)
MiFID	-2.859**	4.874***	-0.421	4.326
	(1.193)	(1.180)	(0.974)	(2.802)
$Lag_1\left(\widehat{HFT}\right)$	1.169***	0.894***	1.293***	1.318***
	(0.109)	(0.108)	(0.085)	(0.097)
$Lag_2\left(\widehat{HFT}\right)$	-0.440***	-0.262***	-0.400***	-0.415***
	(0.104)	(0.097)	(0.084)	(0.100)
N	105	105	105	105
Durbin Watson	1.82	1.86	2.04	2.07
Adjusted R-Squared	0.806	0.726	0.856	0.849
Year Fixed Effects	No	No	No	No

9. Robustness test

To explore the sensitivity of the results in the previous section, we examine a subset of the initial sample chosen to reflect the securities which are most likely to be considered by HFT participants.[22] The securities chosen are from the top traded quintile and are re-balanced on an annual basis. This results in a sample of approximately 570 stocks on the LSE and 220 on Euronext. This sample reflects the most liquid securities in the market and as shown in previous empirical studies such as Brogaard (2010) and Jarnecic and Snape (2010), HFT

[22]Unfortunately, due to the limited number of monthly observations, an analysis of a shorter time period is unlikely to provide stable estimates.

participation exists primarily in large, liquid stocks.[23] As a result we expect to find stronger relationships in this subset than we do when assessing the market in aggregate.

In the interest of brevity, the results of our robustness tests are included in Appendix 1. The findings are qualitatively similar to the results presented in our full sample, with increased levels of HFT associated with lower spreads and integrity violations. Consistent with the findings of previous empirical works, both the significance and magnitude of the impact of HFT on spreads is greater than in our primary sample. In our integrity equations the HFT variable exhibits the same negative correlation with both insider trading and marking the close; however in the restricted sample these coefficients are all highly significant.

10. Conclusion

This study examines the impact of HFT activity on the efficiency and integrity of the Euronext Paris and LSE marketplaces. Whilst a number of academic studies have shown that HFT participation has a positive impact on market efficiency proxied by transaction cost variables, none of these studies have examined the impact of HFT on integrity, nor have any examined the interaction between the multiple dimensions of market quality and HFT.

In order to understand the impact of HFT participation on the LSE and Euronext markets, we have examined a range of key market metrics over an eleven year period, between the years 2001-2011. Our results indicate that there has been significant structural change in the two markets during this period. The average trade size and value have reduced significantly as have average bid-ask spreads. This evidence directly coincides with the introduction of HFT activity, whose participants have been empirically shown to trade in smaller packages. Our results further show significant increases in both the order to trade and cancel to trade ratios, following the 2007 introduction of the MiFID reforms, which are similarly associated with the frequent order to entry and cancellation strategies undertaken by HFT participants.

The reduction in spreads observed over the sample period in this paper is consistent with theoretical predictions which suggest that because HFT participants are better able to update their information sets relative to other market participants, they incur lower adverse selection and inventory holding costs as compared to traditional market makers. We additionally observe from our results significant changes in the millisecond periodicities in the two markets subsequent to 2007. In particular, there are heightened levels of activity observed at a frequency of 100ms for the LSE. Collectively, this evidence indicates a significant and rapid acceleration of HFT activity in both marketplaces.

In order to examine the impact of this increase in HFT activity, our study utilises the cancellation to trade ratio, a proxy commonly used for HFT activity in the academic literature. We utilise relative effective spreads as a proxy for efficiency and construct two measures to proxy for the level of integrity in the market – one proxy for insider trading and the other for end of day manipulation. We integrate these metrics into a three stage simultaneous regression framework, which allows for the potentially endogenous relationships that may exist between each of the dependent variables, whilst also controlling for fixed effects relating to the regulatory environment and time variation.

[23] We would like to thank an anonymous referee for suggesting this line of inquiry.

The results of our simultaneous equations regression framework indicate that the increasing level of HFT participation on the LSE and Euronext-Paris markets have unambiguously increased market efficiency by reducing bid-ask spreads. This result is consistent with the observed decrease in spreads through time and with the theoretical and empirical findings of previous HFT studies. It remains to be determined what the impact of HFT has been on price discovery which we leave for a subsequent work.

Our results regarding integrity suggest that there is a negative relationship between HFT participation and dislocation of the end of day price which we use as a proxy for closing price manipulations. This evidence implies that whilst HFT participants may increase their activity towards the end of day, they do not do so with enough force to dislocate the closing price, quite the opposite. Further, our evidence may indicate that HFT participants actually reduce the incidence of end of day dislocation/manipulation by providing liquidity in times when manipulators may be present, minimising their impact on the market. Our results regarding insider trading are less clear. Whilst it appears that HFT participation is associated with less insider trading on Euronext, it seems to be associated with more insider trading on the LSE. The ambiguous direction of this relationship may be a function of the regulatory design or enforcement differences between the two markets or could be due to other market-specific structural changes not controlled for in the current study.

Overall, our results indicate that the increase in the level of HFT activity has increased efficiency without harming the integrity of the market. Indeed, there is evidence that the integrity of the market has improved as the level of HFT participation has increased. The additional evidence on market integrity, subject to qualifications on the nature of the data available to study (i.e. only publically available data), ameliorates concerns that HFT undermines market quality through reducing market integrity.

References

Aitken, M.J. & deB. Harris, F.H., 2011. Evidence-Based Policy Making for Financial markets: A Fairness and Efficiency Framework for Assessing Market Quality, *The Journal of Trading*, 6(3), pp.22–31.

Brogaard, J.A., 2010. High Frequency Trading and Its Impact on Market Quality. Working Paper, Northwestern University School of Law.

Gonzalo, J., and C. Granger, 1995, Estimation of common long-memory components in co-integrated systems, *Journal of Business and Economic Statistics* 13, 27–35.

Chordia, T., Roll, R.W., Subrahmanyam, A., 2010. Recent Trends in Trading Activity and Market Quality. SSRN eLibrary.

Comerton-Forde, C., Putnins, T.J., 2011. Measuring closing price manipulation. *Journal of Financial Intermediation* 20, 135–158.

Cumming, D., Humphery-Jenner, M., Wu, E., 2011. Exchange Trading Rules, Governance, and Trading Location of Cross-Listed Stocks. SSRN eLibrary.

Cvitanic, J., Kirilenko, A.A., *2010. High Frequency Traders and Asset Prices. SSRN eLibrary.*

Dubow, B., Monteiro, N., 2006*. Measuring Market Cleanliness. Financial Services Authority Occasional Paper 23.*

Frijns, B., Gilbert, A.B., Tourani-Rad, A., 2011*. Do Criminal Sanctions Deter Insider Trading. SSRN eLibrary.*

Gomber, P., Arndt, B., Lutat, M., Uhle, T., 1905. High-Frequency Trading. SSRN eLibrary.

Groth, S., 2011. Does Algorithmic Trading Increase Volatility? Empirical Evidence from the Fully-Electronic Trading Platform Xetra. Wirtschaftinformatik Proceedings 2011.

Harris, L., 2003. Trading and exchanges: market microstructure for practitioners. Oxford University Press.

Hasbrouck, J., 1991. The Summary Informativeness of Stock Trades: An Econometric Analysis. *Review Financial Studies* 4, 571–595.

Hasbrouck, J., 1995. One Security, Many Markets: Determining the Contributions to Price Discovery. *The Journal of Finance* 50, 1175–1199.

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

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

Ito, T., Lyden, S., 2011. Towards an Understanding of High Frequency Trading. Instinet.

Jarnecic, E., Snape, M., 2010. An Analysis of Trades by High Frequency Participants on the London Stock Exchange. Working Paper.

Jovanovic, B., Menkveld, A.J., 2011. Middlemen in Limit-Order Markets. SSRN eLibrary.

Kearns, M., Kulesza, A., Nevmyvaka, Y., 2010a. Empirical Limitations on High Frequency Trading Profitability. SSRN eLibrary.

Kearns, M., Kulesza, A., Nevmyvaka, Y., 2010b. Empirical Limitations on High Frequency Trading Profitability. SSRN eLibrary.

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

Kumar, P., Goldstein, M., Graves, F., 2011. Trading at the Speed of Light: The Impact of High-Frequency Trading on Market Performance, Regulatory Oversight, and Securities Litigation.

Kyle, A.S., 1985. Continuous Auctions and Insider Trading. *Econometrica* 53, 1315–1335.

Lee, C.M.C., Ready, M.J., 1991. Inferring Trade Direction from Intraday Data. *The Journal of Finance* 46, 733–746.

McInish, T., Upson, J., 2012. Strategic Liquidity Supply in a Market with Fast and Slow Traders. Working Paper, University of Texas.

Menkveld, A.J., 2012. High Frequency Trading and the New-Market Makers. SSRN eLibrary.

Prix, J., Loistl, O., Huetl, M., 2007. Algorithmic Trading Patterns in Xetra Orders. *The European Journal of Finance* 13, 717–739.

Venkataraman, K., 2001. Automated Versus Floor Trading: An Analysis of Execution Costs on the Paris and New York Exchanges. *The Journal of Finance* 56, 1445–1485.

Appendix 1: Robustness test regression results

Table A. Top traded quintile - determinants of the effective spread

	Euronext		LSE	
Intercept	5.503	17.105**	0.717	-1.363***
	(3.778)	(7.417)	(0.499)	(0.463)
$\widehat{Number\ of\ Successful\ MTC}$	0.012	-	0.000	-
	(0.008)	-	(0.001)	-
$\widehat{\%\ Information\ Leakage}$	-	3.430***	-	0.367***
	-	(0.610)	-	(0.103)
Ln(Price)	-0.010	-0.267	-0.014	-0.012
	(0.242)	(0.254)	(0.012)	(0.013)
Ln(Traded Volume)	-0.250	0.770**	-0.020	0.056***
	(0.170)	(0.324)	(0.019)	(0.018)
Ln(Volatility)	-0.073	-0.172	0.043***	-0.001
	(0.090)	(0.166)	(0.010)	(0.010)
$\widehat{HFT}$	-0.351***	-0.477***	-0.031**	0.083***
	(0.106)	(0.144)	(0.014)	(0.020)
MiFID	-0.198	-1.584***	-0.038	0.015
	(0.237)	(0.325)	(0.028)	(0.028)
$Lag_1\left(\widehat{Spread}\right)$	0.428***	-0.055	0.822***	0.795***
	(0.100)	(0.110)	(0.047)	(0.085)
N	105	105	105	105
Durbin-Watson	1.51	1.49	2.26	1.91
Adjusted R-Squared	0.462	0.436	0.866	0.747

	Euronext		LSE	
Year Fixed Effects	Yes	Yes	Yes	Yes

Table B. Top traded quintile – determinants of market integrity

	Euronext MTC	Euronext Insider	LSE MTC	LSE Insider
Intercept	-151.893**	-2.723	47.460	0.772
	(74.006)	(2.122)	(73.018)	(1.638)
Ln(Volatility)	5.181***	0.030	1.295	0.013
	(1.716)	(0.049)	(1.979)	(0.042)
Ln(Return)	-56.257***	0.066	-25.874*	-0.283
	(15.109)	(0.162)	(13.809)	(0.276)
Ln(Traded Volume)	8.407**	0.143	-1.685	-0.023
	(3.312)	(0.094)	(3.006)	(0.066)
$\widehat{HFT}$	-6.546***	-0.202***	-6.264**	-0.218***
	(1.100)	(0.035)	(2.503)	(0.027)
Ln($\widehat{Spread}$)	6.722***	0.191**	9.490***	-0.063
	(1.503)	(0.025)	(2.962)	(0.053)
MiFID	6.830***	0.393***	-13.926***	0.133*
	(2.337)	(0.071)	(4.264)	(0.073)
End of Quarter	0.844	0.032**	0.054	0.034***
	(0.575)	(0.012)	(1.414)	(0.010)
$Lag_1\left(\widehat{Marking\ the\ Close}\right)$	0.029	-	0.411***	-
	(0.034)	-	(0.085)	-
$Lag_1\left(\widehat{\%\ Insider\ Trading}\right)$	-	0.004***	-	0.002***
	-	(0.001)	-	(0.001)
N	105	105	105	105

	Euronext MTC	Euronext Insider	LSE MTC	LSE Insider
Durbin-Watson	1.66	1.65	2	2.05
Adjusted R-Squared	0.464	0.155	0.577	0.17
Year Fixed Effects	No	No	No	No

Table C. Top traded quintile - determinants of hft participation

	Euronext		LSE	
Intercept	-54.396**	-27.467	33.076***	2.541
	(22.008)	(18.461)	(9.516)	(27.068)
Ln(Volatility)	2.058***	0.638	0.329	0.590
	(0.518)	(0.398)	(0.250)	(0.693)
Ln(Return)	-20.832***	-3.012	0.003	-3.154
	(5.512)	(2.122)	(1.814)	(4.587)
$\widehat{Number\ of\ Successful\ MTC}$	-0.382***	-	-0.035	-
	(0.051)	-	(0.022)	-
$\widehat{\%\ Information\ Leakage}$	-	-10.047***	-	-20.429***
	-	(1.763)	-	(2.362)
Ln(Spread)	-2.271***	-1.942***	-1.189***	-1.709*
	(0.444)	(0.297)	(0.359)	(0.886)
Ln(Volume)	3.097***	1.535	-1.326***	-0.124
	(1.012)	(0.833)	(0.391)	(1.102)
MiFID	1.252**	3.045***	-1.039	0.205
	(0.575)	(0.501)	(0.716)	(1.161)
$Lag_1\left(\widehat{HFT}\right)$	-0.028	0.151*	0.405***	0.037
	(0.091)	(0.084)	(0.082)	(0.087)
$Lag_2\left(\widehat{HFT}\right)$	0.359***	0.232***	0.221**	0.114
	(0.091)	(0.084)	(0.086)	(0.077)
N	105	105	105	105

	Euronext		LSE	
Durbin Watson	1.79	1.72	1.96	1.98
Adjusted R-Squared	0.596	0.655	0.83	0.569
Year Fixed Effects	No	No	No	No

© Crown copyright 2012
You may re-use this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence. Visit www.nationalarchives.gov.uk/doc/open-government-licence write to the information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gsi.gov.uk

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk

URN: 12/1061