High-frequency trading and execution costs

Amy Kwan

Richard Philip*

Current version: January 13 2015

Abstract

We examine whether high-frequency traders (HFT) increase the transaction costs of slower institutional and retail traders (non-HFT). Using a differences-in-differences test around the introduction of ITCH, a new data feed that decreases the trading latency for HFT, we find that limit order trading costs for non-HFT increase relative to the costs for HFT. We attribute the increase in non-HFT execution costs to more predatory trading by HFT. After the implementation of ITCH, we show that HFT are more successful in front-running non-HFT limit orders, which decreases the execution probability of non-HFT limit orders.

JEL classification: G10, G23

Keywords: High-frequency trading, execution costs, institutional investors

* Corresponding author Email: richard.philip@sydney.edu.au. Amy Kwan and Richard Philip are at the University of Sydney.

1. Introduction

High frequency trading represents one of the most significant changes to market structure in recent years (SEC, 2010). In contrast to slower non-high-frequency traders (non-HFT), HFT respond faster when new information arrives in the market. There are concerns that this speed advantage has created an unequal playing field between short term and longer term (i.e., institutional and retail) traders. While most academic evidence shows that *overall* market quality improves, understanding more clearly how these benefits are shared between fast and slow traders is important.

In this study, we provide evidence that HFT gain from their speed advantage at the expense of slower institutional and retail traders. In April 2012, the Australian Securities Exchange (ASX) introduced ITCH, a new, lower latency data feed, which had the effect of benefitting traders that rely on fast executions. Using the introduction of ITCH as a natural experiment, we use a differences-in-differences framework to compare HFT, institutional and retail execution costs and find an increase in the limit order transaction costs of non-HFT, relative to HFT.

When submitting limit orders, traders face a trade-off between better execution prices, or price improvement, and a risk of non-execution (Foucault, 1999). We develop a new measure for quantifying limit order transaction costs, which captures both the amount of price improvement and the costs associated with non-execution. We find that after the introduction of ITCH, non-HFT limit order trading costs increase, relative to HFT trading costs. Separating limit order transaction costs into its two components, we find that institutional and retail traders receive lower levels of price improvement from their limit orders, which increases overall transaction costs.

We use order-level ASX data to investigate how high-frequency trading strategies impact limit order transaction costs of non-HFT. The data contain broker identifiers so that each incoming message can be attributed to a HFT, institutional or retail broker. Recent studies propose that faster HFT can anticipate and trade ahead of slower non-HFT order flow, which increases the trading costs of non-HFT (see Hirschey (2013), Li (2014), Menkveld (2014)). The introduction of ITCH offers the largest benefits to predatory HFT strategies that rely on speed to front-run, or queue jump, non-HFT order flow. Following the implementation of ITCH, we find a fall in the execution probability of non-HFT limit orders, which reduces the amount of price improvement that non-HFTs receive from their limit orders. This finding provides early evidence that HFTs benefit from lower latencies to more successfully anticipate and trade ahead of non-HFT limit orders. This predatory trading by HFT crowds out the limit order book for other slower traders, thereby increasing their execution shortfall, which leads to higher overall transaction costs.

We quantify front running more directly using a measure of depth imbalance and test whether HFT are more successful in front running non-HFT order flow after the introduction of ITCH. Typically, large order imbalances predict future price movements (Chordia, Roll and Subrahmanyam, 2002, Chordia and Subrahmanyam, 2004). In the context of a limit order book, strategic traders can anticipate future price movements by observing the depth available on the bid and ask side of the limit order book. A large order imbalance on the bid, relative to the ask side of the limit order book provides a noisy signal that buying pressure is likely to increase the future stock price.

A predatory trader in this scenario has several options. First, in anticipation of large future price rises, a strategic market order trader may submit an order to buy at the best ask price when large buying pressure exists on the bid prices.

Second, if the bid-ask spread is wider than one tick, a new limit order can price improve on the existing best bid price, creating a new best bid price. Because of price-time priority, the predatory trader is at the beginning of the queue and will have the highest probability of buying the stock at the best bid price before the price rises.[1]

Third, traders may refrain from trading by strategically withdrawing their limit orders. Large order imbalances on the bid side of the book cause upward pressure on the future stock price and sell limit orders risk being picked off the limit order book. In anticipation of a future price rise, strategic traders will withdraw their sell orders to reduce the risk of trading at a stale price. Each of these predatory trading strategies exacerbate the depth imbalance as volume is removed from the thin side or added to the thick side of the limit order book.[2]

To proxy for queue jumping, we calculate a measure of depth imbalance, defined as the difference between the volumes available on the bid (ask) side and ask (bid) side of the limit order book at the time of buy (sell) order executions.[3] Predatory trading predicts large depth imbalances in the same direction as the order submission; when large depth imbalances build up on the bid side, predatory traders strategically submit buy limit and market orders and withdraw sell limit orders. Similarly, predatory traders are more successful at queue jumping if they submit sell orders when large depth imbalances exist on the ask prices. In support of the front-running hypothesis, we find that depth imbalances immediately before order executions are larger for HFT following the implementation of ITCH. Specifically, HFT are more effective in demanding liquidity on the thin side of the limit order book and providing liquidity when the order book is thick post-ITCH.

Our study extends the growing literature that examines the impact of trading speed on market quality. Early studies on algorithmic trading rely on exogenous events which cause a change in the level of algorithmic trading and find that this trading generally improves market liquidity and price efficiency. Using the introduction of the NYSE automated quote

---

[1] Harris (1999) describes this trading strategy as penny jumping.

[2] Boehmer, Fong and Wu (2014) find that transitory volatility increases with HFT activity. Larger depth imbalances due to predatory HFT trading may contribute to these increases in short-term volatility.

[3] Our results are robust to depth imbalance measures calculated over various levels of the limit order book.

dissemination as an exogenous instrument to measure the effects of algorithmic trading on liquidity, Hendershott, Jones and Menkveld (2011) find that algorithmic trading improves liquidity and quote informativeness. Similarly, in an international study of 42 equity markets around the world, Boehmer, Fong and Wu (2014) find that algorithmic trading improves liquidity and information efficiency but also increases volatility. Malinova, Park and Riordan (2012) use a change in regulatory fees in Canada, which increased the cost for some algorithmic trading strategies, to test the impact of algorithmic trading on retail and institutional trading activity. They conclude that retail and institutional trading costs increase after the reduction in algorithmic trading. Hasbrouck and Saar (2013) proxy for the amount of low latency trading using patterns of order submissions and cancellations and find that low latency trading is associated with higher market quality. Conversely, Eggington, Van Ness and Van Ness (2014) conclude that quote stuffing, a practice associated with HFT trading, can increase trading costs and short term volatility.

Relying on more accurate indicators of HFT trading, which can be considered a subset of algorithmic trading, Brogaard, Hendershott and Riordan (2014) and Hendershott and Riordan (2013) also report improvements in market quality. Using a Nasdaq dataset that identifies a subset of HFT trading, Brogaard, Hendershott and Riordan (2014) find that HFT improve price efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors. Hendershott and Riordan (2013) show that algorithmic traders demand liquidity when bid-ask spreads are narrow and provide liquidity when it is expensive using a sample of algorithmic trades in 30 DAX stocks.

While the focus of these studies is on overall market quality, more recent studies argue that fast traders may benefit at the expense of slower traders. Biais, Foucault and Moinas (2013) show that the presence of fast traders can generate negative externalities by increasing adverse selection costs. Li (2014) models a market in which fast HFT can front-run incoming orders of slower traders, resulting in a transfer of wealth from slower traders to HFT. His model shows that faster HFT can front-run normal-speed traders making markets less liquid and prices less informative. Similarly, Menkveld (2014) argues that HFT may hurt market quality if they aggressively pick off quotes set by other market participants while they may lower adverse-selection costs if acting as market makers. Hirschey (2013) proposes that HFT may increase non-HFT trading costs by anticipating and trading ahead of their order flow. Van Kervel (2014) describes a trading strategy in which HFT, acting as market makers, duplicate their limit orders on several venues to increase execution probabilities before cancelling these orders after observing a trade on one venue. Thus, market-wide measures of depth may overstate the actual liquidity available to investors. However, using more direct measures of institutional execution costs, Brogaard, Hendershott, Hunt and Ysusi (2014) find that institutional trading costs do not change as a result of an increase in HFT. Using a broader set of trader categories, we find that HFT exploit their speed advantage to trade ahead of non-HFT order flow, which increases transaction costs for slower institutional and retail investors.

2. Institutional details

The Australian Securities Exchange (ASX) is the dominant stock exchange for Australian equities, with a 90% market share of on-market traded volume.[4] In 2014, approximately 2050 companies are listed on the ASX with a market capitalization of approximately AUD 1.5 trillion. The ASX operates as a continuous limit order book between 10:00 am and 4:00 pm, which matches orders based on price and time priority. Each stock opens with an opening auction at a random time between 10:00 and 10:10 am depending on the starting letter of their ASX code. Similarly, the closing price is determined via a closing

[4] The remaining 10% of on-market trading takes place on Chi-X Australia, which was launched in October 2011.

price auction that takes place between 4:10 pm and 4:12 pm. While the full order book is available to all market participants, trading on the ASX has been anonymous since the removal of broker identifiers in November 2005.

ASX ITCH

We conduct a difference-in-difference analysis around the introduction of ASX ITCH. Implemented in April 2012, ASX ITCH is the ultra-low latency protocol for accessing ASX market information. ASX ITCH was designed to meet the requirements of speed sensitive traders and increased market information access speeds by up to seven times existing connections (ASX, 2013). Thus, the introduction of ASX ITCH is likely to create larger benefits for HFT, whose strategies may rely on fast response times when new information arrives in the market.

3. Data and sample

We obtain full order book and trade data from the ASX from Thomson Reuters Tick History AusEquities database to examine the impact of the introduction of ITCH on HFT and non-HFT transaction costs. Data from the ASX offer several advantages over other exchanges. First, the data contain broker identifiers for each message so that the initiator of each order and trade can be classified as a HFT, institutional or retail broker. Second, orders and trades are time-stamped to the millisecond allowing a precise reconstruction of the order book. Specifically, we are able to trace every individual order entry, cancellation and amendment, its queue position, and the shape of the limit order book immediately before order submission. We can also determine whether a trade is buyer or seller initiated without relying on trade classification algorithms. Third, in comparison to U.S. equity markets, the

ASX is less fragmented. The ASX operated as a virtual monopoly in Australian equities until the introduction of Chi-X in 2011.[5]

We obtain data for the periods 1 May 2011 – 31 July 2011 and 1 May 2012 – 31 July 2012 for stocks in the S&P/ASX 100 index. The S&P/ASX 100 index contains the 100 largest stocks listed on the ASX by market capitalization. These securities are highly liquid and actively traded among HFT and institutional investors. We compare HFT and non-HFT trading in the post-event period from 1 May 2012 – 31 July 2012 with the pre-event period from 1 May 2011 – 31 July 2011. For the pre-event period, we select the same three months in the preceding calendar year to control for seasonal variations in trading. To ensure that our sample is not contaminated by the opening and closing call auctions, only trades and orders entered between 10:10:00 and 16:00:00 are included. We assume that all outstanding orders remaining in the limit order book at the end of the trading day are cancelled.

For each order, the data contain detailed information on the stock symbol, date and time of order entry, order size and price and broker identifier. Each order has a unique identifier so that subsequent amendments, executions or cancellation can be traced to the original order entry. Broker identifiers are classified into four categories: HFT, institutional, retail and other.[6] Orders originating from institutional and retail brokers are collectively referred to as non-HFT. Our analysis focuses on limit orders that are entered at the best bid or ask prices. We restrict our analysis to this subsample of orders as these orders represent clear intentions to trade.[7] To conceal trading intentions, institutions use algorithms to break up a single large order into multiple smaller orders. Thus, small institutional order flow may have predictable patterns that can be detected by HFT (see Hasbrouck and Saar (2013)). To account for this possibility, for each stock we rank all institutional orders entered at the best

[5] Over our sample period, Chi-X averaged around only 10% of daily trading volume for on-market trades.
[6] Because 'other' brokers can represent either HFT or non-HFT trading, we exclude these orders from the analysis. Broker classifications are based on consultations with industry professionals.
[7] Less aggressive limit orders do not represent a clear intention to trade.

bid or ask prices by size and further separate these orders into large institutional (top quintile) and small institutional (bottom quintile) orders.

Table 1, Panel A reports the summary statistics for the stocks in our sample. The average stock has a *Market capitalization* of 13.52 AUD billion and volume weighted trade price of $11.52. For each stock, approximately 12,000 orders are submitted, which result in 4,010 trade executions. Approximately half of the 12,000 orders entered are subsequently cancelled.

Table 1, Panels B to E report the trading characteristics of HFT, large institutional, small institutional and retail traders, respectively. Institutional orders are classified as large (small) if the order size is in the top (bottom) quintile of all orders entered at the best bid or ask prices for a particular stock. We find that the rate of order cancellations is higher for HFT and large institutional relative to other trader types. On average, HFTs submit 380 orders, of which 247 (65%) are subsequently cancelled. In comparison, retail traders cancel only 15% of their submitted orders, indicating that HFTs monitor their trading strategies more actively than retail traders. Large institutional orders are likely to have a higher execution shortfall compared to both retail and small institutional orders. Consistent with this view, we find that large institutional orders are more frequently cancelled and amended to ensure that these orders do not trade at a stale price.

4. Execution cost measures

4.1 Limit order transaction costs

When submitting limit orders, traders face a trade-off between better execution prices and a risk of non-execution (Foucault, 1999). While market orders allow a trader to execute an order with certainty at prices available in the market, a trader who submits a limit order has the possibility to improve the execution price by buying (selling) at a price below (above) the midpoint of the best bid and ask prices. However, a limit order trader also faces the risk of non-execution if their order is not matched by an incoming market order, in which case the trader will either amend or cancel the original limit order. To capture these two dimensions of limit order submission strategies, we measure total limit order transaction costs (LTC) as the difference between the gains from price improvement and the losses from non-execution.

Price improvement captures the gains to the trader for trading at the best bid price for buys, or the best ask price for sells. However, the amount of price improvement the trader gains is complicated by two factors. First, a limit order can be partially filled if an incoming market order matches only part of the original limit order. In this scenario, the remaining balance of the original order remains in the limit order book until another market order arrives. Second, the limit order trader risk being picked off at a loss if new public information arrives (Foucault, 1999; Liu, 2009). Prices are likely to become stale the longer the order sits in the limit order book. To account for the time dimension, we standardise our measure of price improvement by the waiting time between original order submission and order execution.

To measure the size of the price improvement (*PrcImprove*), we estimate:

$$PrcImprove = q \times \sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{ExecutionPrice_{t+\tau_i} - Mid_t}{Mid_t} \times \frac{1}{\tau_i}$$

where q is a signed indicator variable that takes a value of +1 for orders entered at the best bid price and -1 for orders entered at the best ask price, t is the time the limit order is submitted, Mid_t is the midpoint of the best bid and ask prices at the time of order entry, $Price_{t+\tau_i}$ is the price at the time of execution and $Volume$ is the number of shares entered in the original limit order. To capture the effects of partial executions and waiting time, we weight our measure of $PrcImprove$ by $VolTrade_i$, which is the volume executed for the *i*th partial fill, and scale by τ, the number of events between order entry and order execution.[8]

If an order fails to execute, the limit order trader can either amend the order or delete the order from the limit order book.[9] We measure the costs associated with non-execution (NE) by comparing the bid-ask midpoint at the time of order amendment/cancellation with the bid-ask midpoint at the time of order entry:

$$NE = q \times \frac{VolFail}{Volume} \times \frac{Mid_{t+r} - Mid_t}{Mid_t} \times \frac{1}{r}$$

where $VolFail$ is the number of shares that fail to trade and r is the number of events between order submission and deletion.

Thus, *LTC* is the sum of the benefits of potential price improvement and the costs of non-execution. Formally, this is expressed as:

$$LTC = q \times \left[\sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{ExecutionPrice_{t+\tau_i} - Mid_t}{Mid_t} \times \frac{1}{\tau_i} + \frac{VolFail}{Volume} \times \frac{Mid_{t+r} - Mid_t}{Mid_t} \times \frac{1}{r} \right]$$

Finally, to arrive at daily measure of *LTC*, we weight each order by the order size for each trading day and trader type.

4.2 Depth imbalance

Traditional measures of market depth aggregate book depth across both bid and ask prices (see Hasbrouck and Saar, 2013; Degryse, de Jong and van Kervel, 2011). However, aggregated measures of market depth do not measure the amount of depth available on the side of the limit order book where it is most needed. For example, a trader submitting a buy market order is more concerned about the depth available on the ask side of the limit order book, rather than aggregated depth over both bid and ask prices. Van Kervel (2012) notes that

[8] We measure waiting time in event time as suggested in Hasbrouk and Saar (2013). Each order submission, amendment, cancellation or execution is classified as an event.

[9] We assume that orders remaining in the limit order book at the end of the trading day are cancelled.

aggregated depth over multiple venues can overstate the actual liquidity available to investors as high frequency traders cancel limit orders on the same side of the order book of competing venues after observing a trade on one venue.

To measure the shape of the limit order book at the time of order submission, we calculate depth imbalance (*DepthImbalance*) as the difference between the volume available at the best bid price and the volume available at the best ask price and multiply by an indicator for whether the order is a buy or sell order. When a trader submits a market order, a higher measure of *DepthImbalance* indicates that less liquidity is available on the side of the limit order book where it is demanded. For limit orders, a higher measure of *DepthImbalance* indicates that market makers are providing liquidity on the thick side of the limit order book. Specifically, for each order we calculate:

$$DepthImbalance = q \times \frac{VolBid_{t-\varepsilon} - VolAsk_{t-\varepsilon}}{VolBid_{t-\varepsilon} + VolAsk_{t-\varepsilon}}$$

where $VolBid_{t-\varepsilon}$ ($VolAsk_{t-\varepsilon}$) is the volume available at the best bid (ask) price immediately before order submission at time *t* and q is a signed indicator variable that takes a value of +1 for buy orders and -1 for sell orders. We arrive at a daily measure of *DepthImbalance* by volume weighting over the number of shares corresponding to each submitted order. Higher levels of *DepthImbalance* for a particular trader indicates that they are systematically demanding liquidity from the thin side of the order book, where it is most needed, or supplying liquidity to the thick side of the order book, where it is least needed. Accordingly, *DepthImbalance* can be interpreted as a measure of the front-running activity undertaken by a particular trader type.

Table 2 presents summary statistics for measures of *LTC* and *DepthImbalance*. We find that the average *LTC* is 5.2 basis points for HFT (Table 2, Panel A) and ranges from 0.05 for small institutional orders to 6.9 basis points for retail traders (Table 2, Panels B to Panel D). Separating *LTC* into *PrcImprove* and *NE* components, we find that retail traders receive more price improvement for their limit orders and that large institutional orders have the highest *NE* costs. For *DepthImbalance*, we find that all trader types typically demand liquidity when their own side of the limit order book is thicker and supply liquidity to the thin side of the limit order book, which is consistent with the findings from Ranaldo (2004). Comparing between trader types, it is evident that HFT systematically demand liquidity from the thin side of the order book, relative to other trader types.

5. Empirical results

5.1 Limit order transaction costs

We test whether *LTC* has changed for these trader categories after the introduction of the ITCH data feed in April 2012, which had the effect of reducing latency by up to seven times the previous available levels. Table 3 presents univariate tests on the differences in mean *LTC* between the pre- and post- period. Except for retail limit orders, we find that *LTC* typically falls in the post-ITCH period for all other trader categories after the implementation of ITCH, although the results are not consistently statistically significant. In highly volatile periods, a trader has a higher probability that their limit order will be executed, resulting in a lower *LTC*. Thus, changes in *LTC* over time can be a function of current market conditions and standard event study methodologies may fail to adequately control for these differences. Accordingly, we use a differences-in-differences framework so that the differential impact of ITCH on transaction costs for the treatment group can be compared to that of a control group. In our application, to test whether lower latencies increase non-HFT trading costs, we use HFT as the control group and the non-HFT trader type as the treatment group.

Table 4, Panel A reports the difference-in-differences estimates for total limit order transaction costs. We scale *LTC* by the mean value of *LTC* in the pre-ITCH period. Thus, *LTC* takes a value of 1 for both HFT and non-HFT in the pre-ITCH period, which is reflected in a constant intercept term.[10] We expect the largest increases in the transaction costs of retail and small institutional orders. Given small institutional order flow is likely to be generated from algorithms that could be detected and preyed upon by HFT, large institutional order flow arrival is less predictable. Similarly, retail order flow is less actively monitored and accordingly, is more easily exploited by HFT.[11] Consistent with our predictions, we find that *Non-HFT x Post-ITCH* is consistently positive and significant for small institutional and retail orders. Thus, relative to HFT limit orders, small institutional and retail limit orders are more expensive to trade after the introduction of the faster data feed. For large institutional limit orders, we find that *Non-HFT x Post-ITCH* is positive, but the result is only significant for sell limit orders.

The decision to submit a limit order represents a trade-off between better execution prices (*PrcImprove*) and the cost of non-execution (*NE*). Table 4, Panels B and C decomposes total *LTC* into *PrcImprove* and *NE,* respectively. Because *PrcImprove* has an upper limit of 0, we scale such that *PrcImprove* takes a value of -1 in the pre-period. Thus, a decrease (increase) in *PrcImprove* corresponds to an increase (decrease) in *LTC*. Comparing Table 4, Panels B and C, we find that the increase in non-HFT transaction costs is due to a reduction in *PrcImprove*. In Table 4, Panel B, we find that *Non-HFT x Post-ITCH* is negative and significant across all trader types, indicating that non-HFT limit order traders are receiving less price improvement from their orders post-ITCH, which increases their overall limit order transaction costs. In contrast, Table 4, Panel C shows that *NE* is lower for most trader types post-ITCH, relative to HFT *NE*. Both *PrcImprove* and *NE* are a function of the proportion of the submitted limit order that is successfully executed (*VolTrade/Volume*). In Table 4, Panel D, we examine how *VolTrade/Volume* has changed as a result of the faster data feed. Our results show a significant decrease in the proportion of a non-HFT limit order

[10] Similarly, the coefficient on *Non-HFT* is equal to 0 as *LTC* is scaled to a value of 1 in the pre-ITCH period.
[11] Table 1 shows that *Retail* has the lowest frequencies for order amendments and cancellations.

executing after the implementation of ITCH; across all trader types, we find that *Non-HFT x Post-ITCH* is negative and significant. The increase in non-HFT execution shortfall indicates that HFT are more successful in crowding out non-HFT from the limit order book post-ITCH.

Overall, we find that total *LTC* increases for slower non-HFT, relative to faster HFTs, post implementation of ITCH. Decomposing *LTC* into *PrcImprove* and *NE* components, we find that non-HFT are receiving less price improvement from their orders, which is due to a fall in *VolTrade/Volume*. One potential reason for the decrease in *VolTrade/Volume* is that HFT are able to exploit the faster data feed to anticipate and trade ahead of non-HFT order flow.

5.2 Strategic order placement strategies

The results from the previous section highlight that non-HFT limit order transaction costs, relative to HFT transaction costs, have increased following the implementation of ITCH. Decomposing limit order transaction costs into its individual components, we find that the increase in transaction costs is due to a fall in the execution probabilities of a limit order. This finding is consistent with the notion that faster HFT can anticipate and trade ahead of slower non-HFT order flow, which increases their trading costs (see Hirschey (2013), Li (2014), Menkveld (2014)). In this section, we provide a more thorough examination of how faster data feeds may be differentially affecting HFT and non-HFT trading behaviour.

To measure strategic order placement strategies, we calculate *DepthImbalance*, which compares the amount of liquidity on the bid and ask sides of the limit order book. If traders are successful in anticipating future price movements due to a temporary imbalance in supply and demand, we predict positive *DepthImbalance*, immediately prior to an incoming order.[12] For example, a large depth imbalance on the bid side of the limit order book provides a noisy

[12] We measure *DepthImbalance* immediately before the order is executed, which allows us to account for order submissions that are subsequently cancelled.

signal of a future price rise. A predatory trader will strategically submit buy limit and market orders in anticipation of the future price rise. The reduced latencies after the introduction of ITCH allow speed sensitive traders to react faster to changes in the order book. Traders must react quickly to large imbalances in the order book, either by demanding from the thin side of the order book or by supplying to the thick side of the order book in order to gain queue priority. Thus, we predict an increase in *DepthImbalance* for HFT, relative to non-HFT, after the implementation of ITCH.

Table 5 compares *DepthImbalance* before and after the implementation of ITCH for each trader type. For buy (sell) orders, *DepthImbalance* is calculated by comparing the total depth available for 5 price levels on the bid (ask) side of the limit book with that available on the 5 ask (bid) prices and takes a value between -1 and 1. We average across all observations in the pre (post) period so that we arrive at a single pre (post) value for *DepthImbalance* for each stock. A large *DepthImbalance* signals to the trader that prices are likely to move in the direction of the imbalance. Thus, a trader type with larger values of *DepthImbalance* indicates that the trader is more successful in front-running the order book. For all orders in Table 5, Panel A, we find that HFTs are more strategic with the placement of market and limit orders post-ITCH. We report similar results when separating orders into buy and sell orders. We find that HFTs place buy (sell) market orders when *DepthImbalance* is larger post-ITCH meaning that a larger amount of depth exists on the bid (ask) prices relative to the depth available on the ask (bid) prices. For limit orders, *DepthImbalance* is negative in the pre period and becomes larger in the post period indicating that HFTs are more successful in placing limit orders after the introduction of ITCH. In comparison, we do not find a significant increase in *DepthImbalance* post ITCH for slower, non-HFT, indicating that they are not benefitting from the faster speed.

ITCH is likely to have the greatest benefits for fast traders buying and selling large, actively traded stocks. Thus, we expect to see larger changes in *DepthImbalance* for more active stocks after the implementation of ITCH. In Table 5, Panels B and C, we present results separately for the large and small stock subsamples. Consistent with this view, while we find that *DepthImbalance* increases for HFT limit and market orders across both large and small stock subsamples, the magnitude of the change and the statistical significance of the results are typically larger for the large stock subsample.

Taken together, these results indicate that HFTs are more effective in demanding liquidity on the thin side of the limit order book and providing liquidity when the order book is thick post-ITCH.

Table 6 presents the results for the difference-in-difference regressions of *DepthImbalance* before and after the introduction of ITCH for market orders. Since *DepthImbalance* is bounded by -1 and 1, we scale the measure by subtracting the mean *DepthImbalance* in the pre-ITCH period for each trader type and each stock. Thus, for the difference-in-difference regressions, the intercept and *HFT* take values of 0 to reflect that *DepthImbalance* is scaled to 0 in the pre-period for both HFT and non-HFT. Table 6, Panel A, shows that the interaction term *HFT x Post-ITCH* is positive and significant for all trader types indicating that HFT order placement strategies have improved, relative to non-HFT orders, after the introduction of ITCH. The result is robust to various measures of *DepthImbalance*. Table 6, Panel B repeats the analysis using depth up to 3 levels of the bid and ask prices. Except for small institutional sell orders, we find that *HFT x Post-ITCH* is positive and significant for all other order types.

Table 7 repeats the difference-in-difference analysis of *DepthImbalance* for limit orders. The results are largely consistent with the results from Table 6 analysing market orders. Across all measures of *DepthImbalance,* we find that *HFT x Post-ITCH* is typically positive and significant for each trader type indicating that HFT are more successful in their limit order placement strategies after the implementation of ITCH.

Together, these results confirm our conclusions from the univariate analysis. We find that ITCH provides speed advantages to HFT, enabling them to optimize their order placement strategies. Specifically, they are more successful in anticipating and trading ahead of future price movements based on information contained in the limit order book, which comes at the expense of slower, non-HFT.

6. Conclusion

We use the introduction of ITCH on the ASX, which has a differential impact on fast and slow traders, as a natural experiment to examine the effects of HFT trading on transaction costs. Implemented in April 2012, ASX ITCH increased market information access speeds by up to 7 times existing connections, and thus, offers the greatest benefits to traders that are speed sensitive.

We provide evidence that HFT trading strategies increase the limit order transaction costs of non-HFT. Using a difference-in-difference framework, we find that non-HFT transaction costs increase, relative to HFT transaction costs, after the introduction of ITCH. Decomposing limit order transaction costs into its individual components, we show that the increase in transaction costs is due to an increase in the execution shortfall; the number of shares that successfully execute against incoming market orders decreases significantly post-ITCH.

We find strong evidence that the execution shortfall is be due to more strategic order placement strategies by HFT post-ITCH. Our findings show that they are more successful in front-running non-HFT order flow, which increases non-HFT execution shortfalls. With the benefit of a faster data feed, HFTs are more successful in anticipating future price movements based on the shape of the limit order book. HFTs act on this information by demanding liquidity when the limit order book is thin and supply liquidity when it is thick.

Our findings have implications for the regulation of high frequency trading. For regulation to evolve, we need a clear understanding of how HFT and non-HFT order flow interact in equity markets. Our findings show that HFTs exploit their speed advantages at the expense of slower, long-term investors.

References

Biais, B., Foucault, T. and Moinas, S. (2014) *Equilibrium fast trading*, HEC Paris Research Paper No. 968/2013, unpublished working paper.

Boehmer, E., Fong, K. Y. L. and Wu, J. (2014) *International evidence on algorithmic trading*, Singapore Management University, unpublished working paper.

Brogaard, J., Hendershott, T., Hunt, S. and Ysusi, C. (2014) 'High-Frequency Trading and the Execution Costs of Institutional Investors', *Financial Review*, 49(2), 345-369.

Brogaard, J., Hendershott, T. and Riordan, R. (2014) 'High-Frequency Trading and Price Discovery', *Review of Financial Studies*.

Chordia, T., Roll, R. and Subrahmanyam, A. (2002) 'Order imbalance, liquidity, and market returns', *Journal of Financial Economics*, 65(1), 111-130.

Chordia, T. and Subrahmanyam, A. (2004) 'Order imbalance and individual stock returns: Theory and evidence', *Journal of Financial Economics*, 72(3), 485-518.

Degryse, H., De Jong, F. and Van Kervel, V. (2014) *The impact of dark trading and visible fragmentation on market quality*, TILEC Discussion Paper No. 2011-026, unpublished working paper.

Egginton, J. F., Van Ness, B. F. and Van Ness, R. A. (2014) *Quote stuffing*, University of Mississippi, unpublished working paper.

Foucault, T. (1999) 'Order flow composition and trading costs in a dynamic limit order market', *Journal of Financial Markets*, 2(2), 99-134.

Hasbrouck, J. and Saar, G. (2013) *Low-latency trading*, New York University, unpublished.

Hendershott, T., Jones, C. M. and Menkveld, A. J. (2011) 'Does Algorithmic Trading Improve Liquidity?', *The Journal of Finance*, 66(1), 1-33.

Hendershott, T. and Riordan, R. (2013) 'Algorithmic Trading and the Market for Liquidity', *Journal of Financial and Quantitative Analysis*, 48(04), 1001-1024.

Hirschey, N. (2013) *Do high-frequency traders anticipate buying and selling pressure?*, London Business School, unpublished working paper.

Li, W. (2014) *High frequency trading with speed hierarchies*, University of Maryland, unpublished working paper.

Malinova, K., Park, A. and Riordan, R. (2013) *Do retail traders suffer from high frequency traders?*, University of Toronto, unpublished working paper.

Menkveld, A. J. (2014) 'High-Frequency Traders and Market Structure', *Financial Review*, 49(2), 333-344.

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

Van Kervel, V. (2014) *Market fragmentation and smart order routing technology*, VU University Amsterdam, unpublished working paper.

Table 1.

Summary statistics

Table 1 shows the summary statistics for the trading characteristics of our sample stocks. Panel A reports the average daily price and trade characteristics for the ASX 100. Panels B and D report measures of average daily trading activity for HFT, institutional (orders), institutional (small orders) and retail brokers, respectively. Except for *Market capitalization*, all other variables are measured daily before averaging across the stocks in the sample.

	Obs.	Mean	Std. Dev.	Q1	Median	Q3
Panel A: All stocks						
Market capitalization (bil.)	94	13.52	22.77	2.844	10.00	114.8
Price	94	11.52	13.60	3.045	5.564	14.110
Volume	94	4,128,188	5,234,224	1,140,000	2,687,000	4,756,000
Number of trades	94	4,010	3,376	1,911	2,576	4,523
Number of orders						
Entered	94	11,836	10,196	5,809	7,721	11,970
Amended	94	8,777	9,141	3,816	4,869	8,902
Cancelled	94	6,579	5,881	3,231	4,379	6,499
Trade size	94	1,659	2,247	293.4	790.2	2,048
Panel B: HFT						
Volume	94	227,086	316,129	78,820	151,100	244,300
Number of trades	94	277.4	337.1	79.18	123.9	297.9
Number of orders						
Entered	94	379.4	653.3	98.37	173.4	290.9
Amended	94	311.1	638.4	62.72	105.4	196.6
Cancelled	94	247.0	496.5	63.11	106.6	166.8
Trade size	94	1,899	2,610	325.5	920.7	2,650
Panel C: Institutional (large orders)						
Volume	94	1,803,104	2,319,743	468,700	757,000	2,358,000
Number of trades	94	281.1	248.8	147.20	187.1	294.2
Number of orders						
Entered	94	796.5	699.3	402.00	496.5	844.6
Amended	94	652.5	625.3	312.40	487.0	584.1
Cancelled	94	535.1	489.7	265.30	331.6	573.5
Trade size	94	10,946	18,935	2,010.0	4,525.0	10,470

Table 1- *Continued*

Panel D: Institutional (small orders)						
Volume	94	9,109	6,191	4,387	7,129	12,850
Number of trades	94	616.1	575.2	277.60	392.6	670.6
Number of orders						
Entered	94	787.8	665.8	421.30	538.3	836.9
Amended	94	494.8	539.8	217.60	275.4	457.6
Cancelled	94	338.6	247.2	192.50	251.4	374.2
Trade size	94	21.22	20.97	9.11	14.02	20.95
Panel E: Retail						
Volume	94	268,455	659,915	50,800	91,440	191,100
Number of trades	94	168.0	200.0	59.71	99.06	186.9
Number of orders						
Entered	94	198.2	287.4	52.15	89.35	190.5
Amended	94	94.99	189.8	16.15	39.35	81.24
Cancelled	94	29.77	38.88	8.100	15.66	34.93
Trade size	94	2,376	3,453	437.6	1,120	2,737

Table 2.
Summary statistics

Table 2 reports the average daily transaction costs and depth imbalance for each trader type for stocks in the ASX 100. Panels A to D present the measures for activity for HFT, institutional (large orders), institutional (small orders) and retail brokers, respectively. For each buy and sell order, we measure the limit order transaction cost (LTC) as:

$$LTC = q \times \left[\sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{ExecutionPrice_{t+\tau_i} - Mid_t}{Mid_t} \times \frac{1}{\tau_i} + \frac{VolFail}{Volume} \times \frac{Mid_{t+r} - Mid_t}{Mid_t} \times \frac{1}{r} \right]$$

Where *Volume* is the size of the initial order, $VolTrade_i$ is the volume executed for the *i*th partial fill, $VolFail$ is the number of shares that fail to trade, Mid_t is the midpoint of the best bid and ask prices at the time of order entry *t*, $ExecutionPrice_{t+\tau_i}$ is the price at the time of execution, MID_{t+r} is the midpoint of the best bid and ask prices at the time of order deletion, τ the time between order entry and order execution and *r* is the time between order submission and deletion. The first and second components of *LTC* represent the amount of price improvement (*PrcImprove*) and the cost of non-execution (*NE*), respectively. For each order executed, depth imbalance is measured as:

$$DepthImbalance = q \times \frac{VolBid_{t-\varepsilon} - VolAsk_{t-\varepsilon}}{VolBid_{t-\varepsilon} + VolAsk_{t-\varepsilon}}$$

where $VolBid_{t-\varepsilon}$ ($VolAsk_{t-\varepsilon}$) is the volume available at the best bid (ask) price immediately before order execution at time *t* and q is a signed indicator variable that takes a value of +1 for buy orders and -1 for sell orders. To arrive at a daily measure of *LTC* (*DepthImbalance*), we volume weight total *LTC* (*DepthImbalance*) by order size for each trader type. We calculate *DepthImbalance* separately for limit orders (*Supply*) and market orders (*Demand*).

	Obs.	Mean	Std. Dev.	Q1	Median	Q3
Panel A: HFT						
LTC						
Total LTC	94	5.206	27.246	-0.026	0.006	0.704
PrcImprove	94	-0.436	2.466	-0.142	-0.107	-0.071
NonExec	94	0.156	0.377	0.021	0.031	0.053
Depth imbalance						
Supply	94	-0.020	0.028	-0.037	-0.025	-0.005
Demand	94	0.126	0.047	0.092	0.123	0.163
Panel B: Institutional (large orders)						
LTC						
Total LTC	94	2.817	4.511	0.002	1.305	3.701
PrcImprove	94	-0.477	0.829	-0.307	-0.135	-0.084
NonExec	94	0.585	0.756	0.031	0.265	0.898
Depth imbalance						
Supply	94	-0.029	0.013	-0.036	-0.030	-0.023
Demand	94	0.018	0.016	0.010	0.020	0.028
Panel C: Institutional (small orders)						
LTC						
Total LTC	94	0.051	0.503	-0.040	-0.029	-0.016
PrcImprove	94	-0.153	0.092	-0.202	-0.140	-0.088
NonExec	94	0.072	0.205	0.017	0.026	0.041
Depth imbalance						
Supply	94	-0.045	0.041	-0.065	-0.032	-0.019
Demand	94	0.030	0.020	0.018	0.026	0.037

Table 2 - *Continued*

Panel D: Retail						
LTC						
Total LTC	94	6.885	59.034	-0.089	-0.033	8.026
PrcImprove	94	-3.046	12.729	-0.117	-0.073	-0.040
NonExec	94	0.084	0.219	0.008	0.012	0.018
Depth imbalance						
Supply	94	-0.012	0.023	-0.024	-0.015	-0.002
Demand	94	0.010	0.014	0.004	0.011	0.018

Table 3.
Univariate tests of limit order transaction costs

Table 3 presents univariate tests of transaction costs for each trader type for stocks in the ASX 100. For each buy and sell order, we measure the limit order transaction cost (*LTC*) as:

$$LTC = q \times \left[\sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{ExecutionPrice_{t+\tau_i} - Mid_t}{Mid_t} \times \frac{1}{\tau_i} + \frac{VolFail}{Volume} \times \frac{Mid_{t+r} - Mid_t}{Mid_t} \times \frac{1}{r} \right]$$

where $Volume$ is the size of the initial order, $VolTrade_i$ is the volume executed for the *i*th partial fill, $VolFail$ is the number of shares that fail to trade, Mid_t is the midpoint of the best bid and ask prices at the time of order entry t, $ExecutionPrice_{t+\tau_i}$ is the price at the time of execution, MID_{t+r} is the midpoint of the best bid and ask prices at the time of order deletion, τ the time between order entry and order execution and *r* is the time between order submission and deletion. The first and second components of *LTC* represent the amount of price improvement (*PrcImprove*) and the cost of non-execution (*NE*), respectively. To arrive at a daily measure of *LTC*, we volume weight total *LTC* by order size for each trader type. *Pre* and *Post* show the average daily *LTC* for the periods 1 May 2011 – 31 July 2011 and 1 May 2012 – 31 July 2012, respectively. We conduct a t-test on the differences in mean *LTC* between the *Pre* and *Post* periods. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All limit orders				Buy limit orders				Sell limit orders			
	Pre	Post	Difference		Pre	Post	Difference		Pre	Post	Difference	
HFT												
LTC	7.380	2.478	-4.902		9.041	4.428	-4.614		-0.031	-0.003	0.028	***
PrcImprove	-0.092	-0.414	-0.323		-1.639	-0.606	1.034		-0.061	-0.056	0.005	
NonExec	1.895	2.653	0.758		10.681	5.033	-5.647		0.030	0.053	0.023	***
Institutional (large orders)												
LTC	3.798	1.886	-1.912	**	7.520	5.315	-2.204		-0.020	-0.031	-0.011	**
PrcImprove	-0.646	-0.402	0.244		-1.446	-0.987	0.459		-0.039	-0.053	-0.014	***
NonExec	4.655	4.804	0.149		8.966	6.302	-2.663		0.020	0.022	0.003	
Institutional (small orders)												
LTC	0.159	-0.051	-0.210	**	0.788	-0.023	-0.811		-0.027	-0.050	-0.023	***
PrcImprove	-0.062	-0.075	-0.014	***	-0.064	-0.081	-0.017	***	-0.065	-0.084	-0.019	***
NonExec	0.696	0.048	-0.648		0.852	0.058	-0.794		0.038	0.034	-0.004	

Table 3 - *Continued*

Retail										
LTC	2.896	11.137	8.241	7.672	22.711	15.040	-0.049	-0.039	0.010	**
PrcImprove	-12.100	-3.847	8.254	-20.677	-10.899	9.777	-0.061	-0.057	0.005	
NonExec	12.970	9.077	-3.893	28.348	33.611	5.263	0.012	0.018	0.006	***

Table 4.
Difference-in-difference regressions of limit order transaction costs

Table 4 presents difference-in-difference regressions of transaction costs for each trader type for stocks in the ASX 100. For each buy and sell order, we measure the limit order transaction cost (*LTC*) as:

$$LTC = q \times \left[\sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{ExecutionPrice_{t+\tau_i} - Mid_t}{Mid_t} \times \frac{1}{\tau_i} + \frac{VolFail}{Volume} \times \frac{Mid_{t+r} - Mid_t}{Mid_t} \times \frac{1}{r}\right]$$

where *Volume* is the size of the initial order, $VolTrade_i$ is the volume executed for the *i*th partial fill, $VolFail$ is the number of shares that fail to trade, Mid_t is the midpoint of the best bid and ask prices at the time of order entry *t*, $ExecutionPrice_{t+\tau_i}$ is the price at the time of execution, MID_{t+r} is the midpoint of the best bid and ask prices at the time of order deletion, τ the time between order entry and order execution and *r* is the time between order submission and deletion. The first and second components of *LTC* represent the amount of price improvement (*PrcImprove*) and the cost of non-execution (*NE*), respectively. To arrive at a daily measure of *LTC*, we volume weight total *LTC* by order size for each trader type for each trading day and then average across the pre-ITCH period (1 May 2011 – 31 July 2011) and post-ITCH period (1 May 2012 – 31 July 2012). We scale the measures by dividing by the pre-ITCH mean. *Non-HFT* is an indicator variable equal to 1 for the trader category in the column heading and equal to 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs after the introduction of ITCH. Panels A shows the results for total *LTC* and Panels B to D present results for the components of *LTC*: *PrcImprove*, *NE* and *Voltrade/Volume*. Standard errors are reported in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All			Buys			Sells		
	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail
Panel A: Total LTC									
Non-HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.1325)	(0.1593)	(0.1361)	(0.1322)	(0.1504)	(0.1429)	(0.1370)	(0.1532)	(0.1231)
Post - ITCH	-0.7927 ***	-0.7927 ***	-0.7927 ***	-0.8149 ***	-0.8149 ***	-0.8149 ***	-0.7141 ***	-0.7141 ***	-0.7141 ***
	(0.1451)	(0.1745)	(0.1491)	(0.1448)	(0.1647)	(0.1565)	(0.1452)	(0.1624)	(0.1304)
Non-HFT × Post-ITCH	1.3782 ***	0.2724	0.3851 *	1.3781 ***	0.3535	0.3712 *	1.2863 ***	0.7497 ***	0.5495 ***
	(0.1980)	(0.2428)	(0.2070)	(0.1976)	(0.2315)	(0.2156)	(0.2006)	(0.2247)	(0.1796)
Constant	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***
	(0.0937)	(0.1126)	(0.0962)	(0.0935)	(0.1063)	(0.1010)	(0.0968)	(0.1083)	(0.0870)
Obs.	332	320	321	332	315	326	342	341	345
Adj. R-square	0.2036	0.0802	0.0996	0.2045	0.0900	0.0962	0.1783	0.0694	0.094

Table 4 - *Continued*

Panel B: PrcImprove									
Non-HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0717)	(0.0776)	(0.0759)	(0.0752)	(0.0806)	(0.0845)	(0.0683)	(0.0735)	(0.0712)
Post - ITCH	0.6792 ***	0.6792 ***	0.6792 ***	0.7156 ***	0.7156 ***	0.7156 ***	0.5575 ***	0.5575 ***	0.5575 ***
	(0.0717)	(0.0776)	(0.0759)	(0.0752)	(0.0806)	(0.0845)	(0.0683)	(0.0735)	(0.0712)
Non-HFT × Post-ITCH	-0.5045 ***	-0.2830 **	-0.8096 ***	-0.4958 ***	-0.3112 ***	-0.7480 ***	-0.3192 ***	-0.1845 *	-0.6346 ***
	(0.1014)	(0.1097)	(0.1075)	(0.1063)	(0.1140)	(0.1196)	(0.0966)	(0.1040)	(0.1008)
Constant	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***
	(0.0507)	(0.0549)	(0.0537)	(0.0531)	(0.0570)	(0.0597)	(0.0483)	(0.0520)	(0.0504)
Obs.	364	364	363	364	364	363	364	364	364
Adj. R-square	0.2443	0.2265	0.2741	0.2453	0.2302	0.2298	0.1927	0.1869	0.2144
Panel C: NE									
Non-HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.1052)	(0.1363)	(0.1445)	(0.1421)	(0.1655)	(0.1755)	(0.0937)	(0.1414)	(0.1171)
Post - ITCH	0.4669 ***	0.4669 ***	0.4669 ***	0.3152 **	0.3152 *	0.3152 *	0.5021 ***	0.5021 ***	0.5021 ***
	(0.1124)	(0.1456)	(0.1543)	(0.1489)	(0.1735)	(0.1840)	(0.0956)	(0.1443)	(0.1195)
Non-HFT × Post-ITCH	-0.5014 ***	-0.9348 ***	-1.5135 ***	-0.3139	-0.6645 ***	-1.1770 ***	-0.4403 ***	-0.4167 **	0.1447
	(0.1542)	(0.2044)	(0.2210)	(0.2064)	(0.2467)	(0.2626)	(0.1338)	(0.2044)	(0.1680)
Constant	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***
	(0.0744)	(0.0964)	(0.1021)	(0.1005)	(0.1170)	(0.1241)	(0.0662)	(0.1000)	(0.0828)
Obs.	343	328	318	347	331	329	357	349	354
Adj. R-square	0.0623	0.0959	0.21	0.0097	0.03	0.1014	0.0893	0.0368	0.1142
Panel D: Voltrade/Volume									
Non-HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0656)	(0.0663)	(0.0663)	(0.0680)	(0.0683)	(0.0686)	(0.0640)	(0.0645)	(0.0643)
Post - ITCH	0.4684 ***	0.4684 ***	0.4684 ***	0.4931 ***	0.4931 ***	0.4931 ***	0.3615 ***	0.3615 ***	0.3615 ***
	(0.0658)	(0.0664)	(0.0665)	(0.0680)	(0.0683)	(0.0686)	(0.0642)	(0.0647)	(0.0645)
Non-HFT × Post-ITCH	-0.5569 ***	-0.3350 ***	-0.6443 ***	-0.5522 ***	-0.3763 ***	-0.6903 ***	-0.4004 ***	-0.2153 **	-0.5814 ***
	(0.0930)	(0.0938)	(0.0939)	(0.0962)	(0.0966)	(0.0970)	(0.0906)	(0.0913)	(0.0911)
Constant	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***	1.0000 ***
	(0.0464)	(0.0469)	(0.0469)	(0.0481)	(0.0483)	(0.0485)	(0.0453)	(0.0456)	(0.0455)
Obs.	363	363	363	364	364	364	363	363	363
Adj. R-square	0.1902	0.1489	0.2173	0.1866	0.1563	0.2287	0.118	0.0969	0.1819

Table 5.

Univariate tests of depth imbalance

Table 5 presents univariate tests of depth imbalance for each trader type for stocks in the ASX 100. For each order executed, depth imbalance is measured as:

$$DepthImbalance = q \times \frac{VolBid_{t-\varepsilon} - VolAsk_{t-\varepsilon}}{VolBid_{t-\varepsilon} + VolAsk_{t-\varepsilon}}$$

where $VolBid_{t-\varepsilon}$ ($VolAsk_{t-\varepsilon}$) is the volume available at the best bid (ask) price immediately before order execution at time t and q is a signed indicator variable that takes a value of +1 for buy orders and -1 for sell orders. To arrive at a daily measure of *DepthImbalance*, we volume weight total *LTC* (*DepthImbalance*) by order size for each trader type. We calculate *DepthImbalance* separately for market orders (*Demand*) and limit orders (*Supply*). Panels A to C show the results for all stocks, large stocks (above median) and small stocks (below median). *Pre* and *Post* show the average daily *DepthImbalance* for the periods 1 May 2011 – 31 July 2011 and 1 May 2012 – 31 July 2012, respectively. We conduct a t-test on the differences in mean *DepthImbalance* between the *Pre* and *Post* periods. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All orders				Buy orders				Sell orders			
	Pre	Post	Difference		Pre	Post	Difference		Pre	Post	Difference	
Panel A: All stocks												
HFT												
Demand	0.110	0.143	0.033	***	0.120	0.152	0.032	***	0.100	0.134	0.034	***
Supply	-0.036	-0.005	0.030	***	-0.021	0.003	0.024	***	-0.050	-0.013	0.037	***
Institutional (large orders)												
Demand	0.023	0.012	-0.011	***	0.040	0.020	-0.020	***	0.006	0.004	-0.002	
Supply	-0.027	-0.030	-0.003		-0.009	-0.021	-0.011	*	-0.045	-0.039	0.006	
Institutional (small orders)												
Demand	0.029	0.031	0.002		0.046	0.038	-0.008		0.012	0.025	0.012	
Supply	-0.043	-0.048	-0.006		-0.026	-0.041	-0.015		-0.060	-0.056	0.003	
Retail												
Demand	0.012	0.009	-0.003		0.031	0.018	-0.013	*	-0.008	0.000	0.007	
Supply	-0.013	-0.010	0.003		0.006	-0.004	-0.009		-0.032	-0.016	0.015	**

Table 5 - *Continued*

Panel B: Large stocks												
HFT												
Demand	0.094	0.136	0.042	***	0.095	0.134	0.039	***	0.093	0.138	0.045	***
Supply	-0.022	0.012	0.034	***	-0.014	0.010	0.024	***	-0.031	0.014	0.044	***
Institutional (large orders)												
Demand	0.024	0.011	-0.012	***	0.031	0.008	-0.023	***	0.016	0.015	-0.001	
Supply	-0.027	-0.028	-0.001		-0.019	-0.031	-0.012	*	-0.035	-0.025	0.010	
Institutional (small orders)												
Demand	0.034	0.032	-0.001		0.039	0.026	-0.013	*	0.028	0.038	0.010	
Supply	-0.049	-0.055	-0.006		-0.042	-0.060	-0.017		-0.055	-0.049	0.006	
Retail												
Demand	0.012	0.005	-0.007	***	0.023	0.001	-0.021	***	0.002	0.008	0.007	
Supply	-0.019	-0.015	0.003		-0.006	-0.025	-0.020	***	-0.032	-0.005	0.027	***
Panel C: Small stocks												
HFT												
Demand	0.125	0.150	0.025	**	0.143	0.170	0.027	*	0.106	0.129	0.023	
Supply	-0.048	-0.021	0.027	***	-0.028	-0.004	0.024	**	-0.068	-0.038	0.030	***
Institutional (large orders)												
Demand	0.023	0.013	-0.010	***	0.049	0.032	-0.017		-0.003	-0.005	-0.002	
Supply	-0.028	-0.032	-0.004		0.000	-0.012	-0.011		-0.055	-0.052	0.003	
Institutional (small orders)												
Demand	0.025	0.030	0.006		0.052	0.049	-0.003		-0.002	0.012	0.014	
Supply	-0.037	-0.043	-0.005		-0.011	-0.022	-0.012		-0.064	-0.063	0.001	
Retail												
Demand	0.011	0.012	0.001		0.038	0.033	-0.005		-0.016	-0.009	0.008	
Supply	-0.008	-0.005	0.003		0.016	0.017	0.001		-0.031	-0.027	0.005	

Table 6.
Difference-in-difference regressions of depth imbalance (market orders)

Table 6 presents difference-in-difference regressions of depth imbalance for each trader type for stocks in the ASX 100. For each market order, depth imbalance is measured as:

$$DepthImbalance = q \times \frac{VolBid_{t-\varepsilon} - VolAsk_{t-\varepsilon}}{VolBid_{t-\varepsilon} + VolAsk_{t-\varepsilon}}$$

where $VolBid_{t-\varepsilon}$ ($VolAsk_{t-\varepsilon}$) is the volume available at the best bid (ask) price immediately before order execution at time *t* and q is a signed indicator variable that takes a value of +1 for buy orders and -1 for sell orders. To arrive at a daily measure of *DepthImbalance*, we volume weight *DepthImbalance* by order size for each trader type for each trading day and then average across the pre-ITCH period (1 May 2011 – 31 July 2011) and post-ITCH period (1 May 2012 – 31 July 2012). We scale the measures by subtracting the pre-ITCH mean. *HFT* is an indicator variable equal to 1 for HFT and 0 for the trader category in the column heading. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs after the introduction of ITCH. Panels A and B show the results for *DepthImbalance* measured over 5 and 3 price levels in the limit order book, respectively. In Panel C, *DepthImbalance* is measured over 4 price levels, after excluding the depth available at the best bid and ask prices. Standard errors are reported in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All			Buy market orders			Sell market orders		
	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail
Panel A: Depth (5 levels)									
HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0039)	(0.0040)	(0.0040)	(0.0071)	(0.0071)	(0.0072)	(0.0068)	(0.0068)	(0.0069)
Post - ITCH	0.0026	-0.0114 ***	-0.0024	-0.0087	-0.0216 ***	-0.0132 *	0.0140 **	-0.0013	0.0079
	(0.0039)	(0.0040)	(0.0040)	(0.0071)	(0.0071)	(0.0072)	(0.0068)	(0.0068)	(0.0069)
HFT × Post-ITCH	0.0268 ***	0.0408 ***	0.0318 ***	0.0357 ***	0.0485 ***	0.0402 ***	0.0184 *	0.0337 ***	0.0246 **
	(0.0055)	(0.0056)	(0.0057)	(0.0100)	(0.0101)	(0.0102)	(0.0096)	(0.0096)	(0.0098)
Constant	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0028)	(0.0028)	(0.0028)	(0.0050)	(0.0050)	(0.0051)	(0.0048)	(0.0048)	(0.0049)
Obs.	364	364	364	364	364	364	364	364	364
Adj. R-square	0.1769	0.2362	0.1851	0.0662	0.1071	0.0758	0.0716	0.0821	0.0675

Table 6 - *Continued*

Panel B: Depth (3 levels)									
HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0045)	(0.0044)	(0.0043)	(0.0057)	(0.0059)	(0.0060)	(0.0063)	(0.0061)	(0.0062)
Post - ITCH	0.0103 **	-0.0106 **	-0.0025	0.0063	-0.0152 ***	-0.0062	0.0144 **	-0.006	0.0016
	(0.0045)	(0.0044)	(0.0043)	(0.0057)	(0.0059)	(0.0060)	(0.0063)	(0.0061)	(0.0062)
HFT × Post-ITCH	0.0131 **	0.0340 ***	0.0259 ***	0.0178 **	0.0393 ***	0.0303 ***	0.0085	0.0289 ***	0.0213 **
	(0.0063)	(0.0063)	(0.0061)	(0.0081)	(0.0083)	(0.0085)	(0.0089)	(0.0086)	(0.0087)
Constant	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0032)	(0.0031)	(0.0031)	(0.0041)	(0.0041)	(0.0043)	(0.0044)	(0.0043)	(0.0044)
Obs.	364	364	364	364	364	364	364	364	364
Adj. R-square	0.0862	0.1409	0.1085	0.0541	0.1063	0.0683	0.0431	0.0605	0.0439

Table 7.

Difference-in-difference regressions of depth imbalance (limit orders)

Table 7 presents difference-in-difference regressions of depth imbalance for each trader type for stocks in the ASX 100. For each limit order executed, depth imbalance is measured as:

$$DepthImbalance = q \times \frac{VolBid_{t-\varepsilon} - VolAsk_{t-\varepsilon}}{VolBid_{t-\varepsilon} + VolAsk_{t-\varepsilon}}$$

where $VolBid_{t-\varepsilon}$ ($VolAsk_{t-\varepsilon}$) is the volume available at the best bid (ask) price immediately before order execution at time t and q is a signed indicator variable that takes a value of +1 for buy orders and -1 for sell orders. To arrive at a daily measure of *DepthImbalance*, we volume weight *DepthImbalance* by order size for each trader type for each trading day and then average across the pre-ITCH period (1 May 2011 – 31 July 2011) and post-ITCH period (1 May 2012 – 31 July 2012). We scale the measures by subtracting the pre-ITCH mean. *HFT* is an indicator variable equal to 1 for HFT and 0 for the trader category in the column heading. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs after the introduction of ITCH. Panels A and B show the results for *DepthImbalance* measured over 5 and 3 price levels in the limit order book, respectively. In Panel C, *DepthImbalance* is measured over 4 price levels, after excluding the depth available at the best bid and ask prices. Standard errors are reported in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All			Buy limit orders			Sell limit orders		
	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail	Insto. (small)	Insto. (large)	Retail
Panel A: DepthImbalance (5 levels)									
HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0041)	(0.0039)	(0.0045)	(0.0076)	(0.0076)	(0.0085)	(0.0072)	(0.0072)	(0.0080)
Post - ITCH	-0.0064	-0.0026	0.0041	-0.0171 **	-0.0123	-0.0143 *	0.0041	0.007	0.0191 **
	(0.0041)	(0.0039)	(0.0045)	(0.0076)	(0.0076)	(0.0085)	(0.0072)	(0.0072)	(0.0080)
HFT × Post-ITCH	0.0474 ***	0.0435 ***	0.0368 ***	0.0474 ***	0.0425 ***	0.0445 ***	0.0470 ***	0.0441 ***	0.0320 ***
	(0.0058)	(0.0056)	(0.0064)	(0.0108)	(0.0107)	(0.0121)	(0.0102)	(0.0102)	(0.0113)
Constant	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0029)	(0.0028)	(0.0032)	(0.0054)	(0.0053)	(0.0060)	(0.0051)	(0.0051)	(0.0057)
Obs.	364	364	364	364	364	364	364	364	364
Adj. R-square	0.3168	0.3154	0.2349	0.0926	0.0798	0.0668	0.1582	0.1563	0.1245

Table 7 - *Continued*

Panel B: DepthImbalance (3 levels)									
HFT	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0048)	(0.0046)	(0.0052)	(0.0069)	(0.0067)	(0.0077)	(0.0062)	(0.0062)	(0.0074)
Post - ITCH	-0.0121 **	-0.0011	-0.0004	-0.0150 **	-0.0048	-0.0159 **	-0.0092	0.0027	0.0128 *
	(0.0048)	(0.0046)	(0.0052)	(0.0069)	(0.0067)	(0.0077)	(0.0062)	(0.0062)	(0.0074)
HFT × Post-ITCH	0.0303 ***	0.0193 ***	0.0185 **	0.0265 ***	0.0163 *	0.0274 **	0.0339 ***	0.0220 **	0.0118
	(0.0068)	(0.0065)	(0.0073)	(0.0098)	(0.0094)	(0.0109)	(0.0088)	(0.0088)	(0.0105)
Constant	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.0034)	(0.0032)	(0.0037)	(0.0049)	(0.0047)	(0.0054)	(0.0044)	(0.0044)	(0.0052)
Obs.	364	364	364	364	364	364	364	364	364
Adj. R-square	0.094	0.0566	0.0416	0.0317	0.0094	0.0266	0.0757	0.0506	0.033