High-frequency trading and execution costs

Amy Kwan*

Richard Philip

Current version: June 12 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 a new data feed that decreases HFT latencies, we find that non-HFT limit order trading costs increase relative to the costs for HFT. We attribute the increase in non-HFT execution costs to more predatory HFT. After the reduction in trading latencies, we show that HFT are more successful at trading ahead of non-HFT limit orders. The execution probability of non-HFT limit orders falls, thereby increasing the costs of their limit order strategies.

JEL classification: G10, G23

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

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

1.0 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.[1] 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 the academic evidence shows that overall market quality improves (see Jones, 2013; Brogaard, Hendershott, and Riordan, 2013; Carrion, 2013; Menkveld, 2013), recently some researchers have started to reveal a more predatory role of HFT (van Kervel, 2014; Hoffman, 2014; O'Hara, 2015). Given the controversy surrounding HFT and the considerable amounts of resources invested by some market participants to gain speed advantages of a fraction of a second, it is important to more clearly understand how these market-wide benefits are shared between fast and slow traders.

We examine the impact of HFT trading strategies on non-HFT transaction costs using the introduction of a new, lower latency data feed as a natural experiment. In April 2012, the Australian Securities Exchange (ASX) introduced ITCH, which reduced trading latencies for traders relying on fast execution strategies. Using a differences-in-differences framework around the introduction of ITCH and a direct measure of limit order transactions costs, which captures the two dimensions of limit orders—the benefits of price improvement and the costs of non-execution—we analyze the impact of faster trading speeds on HFT, institutional and retail execution costs. Finally, to investigate the channels through which HFT trading strategies impact non-HFT transaction costs, we adopt a proxy for predatory HFT activity, which characterizes the shape of the order book at the time of the trade or order cancellation.

[1] While high frequency trading can be used to describe a large set of trading activities and behaviours (O'Hara, 2015), we focus on the trading activities of pure proprietary HFT firms, who have no commitment to provide liquidity continuously.

Our analysis of the Australian market overcomes many of the limitations of U.S. studies outlined in O'Hara (2015).[2] First, the granularity of the Australian data allows us to estimate the costs of limit order strategies, which have become highly prevalent in the high frequency trading environment (O'Hara, 2015).[3] When submitting limit orders, traders face a trade-off between better execution prices, or price improvement, and a risk of non-execution (Foucault, 1999). The data allow us to reconstruct the full order book by tracking each order from the time of submission to the time when the order is executed, amended or cancelled. Thus, we can capture both the amount of price improvement a trader receives from the proportion of the order that is successfully executed, and the costs the trader must incur when the order is subsequently amended or cancelled at a less favorable price. Our main analysis focuses on limit orders submitted to the best bid or ask prices to ensure that our transaction cost measure includes only orders entered by traders with a genuine intention to trade.

Second, we can directly observe HFT strategies by capturing the precise shape of the order book at the time of a trade or order cancellation. Our measure of predatory HFT is based on the framework of predatory trading proposed by Brunnermeier and Pedersen (2005). In their model, predators initially trade in the same direction as the large trader, thereby increasing the large trader's transaction costs. In a high frequency world, traders with no fundamental information can become informed by predicting market movements better than other traders (O'Hara, 2015). Further, O'Hara (2015) notes that 'high frequency traders want to be at the front of the queue when an attractive order arrives' (p. 3). Since large order imbalances predict future price movements (Chordia, Roll and Subrahmanyam, 2002; Chordia and Subrahmanyam, 2004),

[2] Comerton-Forde and Putnins (2015) outline the benefits of Australian data on studying the effects of dark trading.
[3] In Australian markets, each order and trade is time-stamped to the millisecond and trading venues do not have differing latencies in their trade reporting mechanisms. The data also contain broker identifiers so that each incoming message can be attributed to a HFT, institutional or retail broker.

strategic traders can forecast future prices by observing the depth available on the bid and ask side of the limit order book. If there is a large order imbalance on the bid, relative to the ask side of the limit order book, the strategic trader observes a noisy signal that buying pressure is likely to increase the future stock price, and will want to buy at the front of the limit order queue. On the other hand, a strategic trader will cancel any existing sell orders to avoid being picked off the limit order book. To proxy for predatory HFT, we use a measure of depth imbalance, which is based on the difference between the volumes available on the bid and ask sides of the order book. In anticipation of future price movements, a strategic trader will therefore trade in the same direction of a large depth imbalance and will cancel their orders if the large imbalance exists on the opposite side of the order book. Because improved trading speeds enable fast traders to access and process market information more quickly than slower traders, we use the adoption of ITCH to test whether HFT have become more predatory with their speed advantage.

Finally, compared to the U.S. markets, where trading is fragmented across 11 lit equity markets and 50 or more other trading venues, trading in the Australian market remains largely consolidated during our sample period. Since there is only a limited choice of trading venues for HFT, the consolidated order book allows us to directly investigate pure HFT strategies within the same market without the complications of cross-market and latency arbitrage between trading venues. For these reasons, the ASX provides an ideal setting to investigate the impact of HFT activity on market quality.

We find that the costs of implementing limit order strategies for non-HFT rise, relative to HFT, after the introduction of ITCH. Decomposing limit order transaction costs into its two components—the benefits of price improvement and non-execution costs—we attribute the increase in costs to higher non-execution costs. Thus, when fast traders can access and process information at lower latencies, slower traders face lower execution probabilities and higher risks of adverse selection for their limit orders. In our second line of analysis, we find evidence that HFT activity becomes more predatory when their trading latencies decrease. Specifically, HFT are more successful in trading ahead of large favorable depth imbalances and are more successful at cancelling their orders when depth imbalances arise in the opposite direction. More predatory HFT activity is one mechanism through which HFT increases the costs of non-HFT limit order strategies.

Our study contributes to the literature in three ways. First, we contribute to the ongoing debate about the benefits and concerns related to HFT. Using traditional measures of market quality, earlier studies find that HFT market making reduces spreads and enhances informational efficiency (see Jones, 2012; Brogaard, Hendershott and Riordan, 2013; Carrion, 2013; Menkveld, 2013; Hagstromer and Norden, 2013). However, in contrast to more traditional market makers, HFT market makers are not required to provide liquidity continuously, leading to concerns that HFT can induce market instability (O'Hara, 2015). In support of the destabilizing view of HFT market making, our findings show that HFT supply liquidity on the thick side of the order book, where it is not required, and demand liquidity from the thin side of book, where it is most needed.

More recent studies investigate how these benefits are split between faster and slower traders. 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. Malinova, Park and Riordan (2014) show that a change to regulatory fees in Canada reduced the returns to retail limit orders while the intraday returns to institutional market orders increased. They attribute the decrease in retail traders' intraday returns to less algorithmic trading activity. We extend this literature by examining the specific mechanism through which HFT activity affects non-HFT transaction costs.

Second, we contribute to the growing literature that models the predatory nature of HFT. 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. Hoffman (2014) models a limit order market in which fast traders revise their quotes quickly after the arrival of news to reduce the risk of being picked off the order book, thereby increasing trading volumes. On the other hand, slower traders face higher adverse selection risks and strategically submit limit orders with lower execution probability, which reduces trading activity. Our study provides empirical support for many predictions of these theoretical models. Specifically, our results show that predatory HFT behaviour is an important channel through which HFT activity increases non-HFT transaction costs.

Finally, we contribute to the literature that addresses the changing nature of market microstructure. O'Hara (2015) argues that traditional empirical techniques, such as realized spreads and permanent and transitory price effects, may no longer be appropriate in a high frequency environment. Similarly, many of the traditional measures of market quality were developed from the perspective of a market order trader hitting the specialist quotes or crossing the spread in the limit order book. Today, dynamic limit order strategies dominate the market and existing market quality measures may not adequately capture the costs of these execution strategies. We extend this avenue of research by analysing a measure of market quality, which captures the direct cost of using a limit order.[4] Additionally, market-wide measures of depth may overstate the actual liquidity available to investors. 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. Our measure of depth imbalance captures both the amount of liquidity and the directionality of the liquidity, which overall measures of total depth do not capture. Using more precise measures of market quality, we show that faster market participants use their speed advantage to extract rents from slower traders, which have implications for the fairness of today's equity markets.

2.0 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.[5] In 2014, approximately 2,050 companies are listed on the ASX with a total 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, matching 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 price auction that takes place between 4:10 pm and 4:12 pm. While trading on the ASX has been anonymous since the removal of real-

[4] Brogaard et al. (2012) measure institutional trading costs by comparing execution costs to the volume-weighted average price of the trading day. Similarly, Malinova, Park and Riordan (2014) compare buy and sell dollar volumes to the closing price of the day. The granularity of our dataset allows us to determine the cost of each individual limit order.

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

time broker identifiers in November 2005, the full order book with broker identification is available to all market participants on a *t*+3 basis. From these broker identifiers, we can determine whether an order or trade originates from an HFT, institutional or retail brokerage firm.

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, which can be accessed by all market participants for a monthly fee. 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 rely on fast response times when new information arrives in the market.

3.0 Data and sample

We obtain full order book and trade data for stocks in the S&P/ASX 100 index from the AusEquities database provided by the Securities Industry Research Centre of Asia Pacific.[6] The securities contained in our dataset are highly liquid and actively traded among HFT and institutional investors. To allow time for HFT firms to adapt to the faster ITCH data feed, we analyze data for the periods 1 January 2012 – 31 March 2012 (pre-ITCH) and 1 May 2012 – 31 July 2012 (post-ITCH). We include only trades and orders entered between 10:10:00 and 16:00:00 to ensure that our sample is not contaminated by the opening and closing call auctions. We assume that all outstanding orders remaining in the limit order book at the end of the trading day are cancelled.

[6] The S&P/ASX 100 index contains the 100 largest stocks listed on the ASX by market capitalization.

Data from the ASX offer several advantages over other exchanges. For each order, the data contain detailed information on the stock symbol, date and time of order entry to the millisecond level, order size and price and an identifier for the submitting broker. In our dataset, broker identifiers are classified into four categories: proprietary HFT firms, institutional, retail and other.[7] We refer to orders originating from institutional and retail brokers collectively as non-HFT. Additionally, each order has a unique identifier such that subsequent amendments, executions or cancellation can be traced to the original order entry, allowing for a full reconstruction of the limit order book. We rely on the granularity of the data to compute direct measures of limit order costs as well as our depth imbalance proxy for predatory trading. Our analysis focuses on limit orders that are entered at the best bid or ask prices as this subsample represent orders submitted by brokers with clear intentions to trade.[8] Furthermore, because the data contains information on the broker submitting the initial order, we do not have to rely on trade classification algorithms, such as Lee and Ready (1991), to determine whether a trade is buyer or seller initiated.[9] Finally, in comparison to U.S. and European equity markets, the ASX is less fragmented, operating as a virtual monopoly in Australian equities until the introduction of Chi-X in 2011.[10]

Table 1, Panel A reports the summary statistics for the 94 stocks, which appear in the S&P/ASX 100 index for both the pre- and post- sample periods. The average stock has a *Market capitalization* of 13.52 AUD billion and volume weighted trade price of $11.02. Approximately 4 million shares trade a day with an average trade size of 1,584 shares. The minimum pricing

---

[7] 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.
[8] Less aggressive limit orders do not represent a clear intention to trade.
[9] Easley, Lopez de Prado and O'Hara (2012) and O'Hara (2015) discuss the potential problems with trade classification algorithms in a high frequency environment.
[10] Over our sample period, Chi-X averaged around only 10% of daily trading volume for on-market trades.

increment on the ASX is $0.01 for stocks priced above $2.00. Given an average *QSpread* of 1.02cents, many stocks in the sample are likely to be spread constrained.

Table 1, Panel B represents the trade characteristics for our three broker categories. *HFT*, which contains only proprietary HFT order flow, trades on average 123,016 shares per stock. Based on volume traded, *Institutional* is the most active broker category with approximately 2.5million shares traded per stock each day. However, some of this volume may be due to HFT strategies operated through institutional brokers, which we cannot classify through broker identifiers alone. In Section 4.3, we conduct robustness tests to ensure our results are not driven by broker misclassifications. The average trade size for *Retail* is much higher than for *HFT* and *Institutional*. This may be because retail orders, unlike institutional orders, are not sliced into smaller trade sizes by algorithms. Average trade sizes for *HFT* and *Institutional* are larger than those reported in U.S. studies, which is due to the lower average prices for stocks listed on the ASX.

4.0 Limit order transaction costs

4.1 Limit order transaction costs measure

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 displayed in the limit order book, 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 (*PrcImprove*) captures the gains to the trader for trading at the best bid price for buys, or the best ask price for sells. 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. To capture the total proportion of the order that receives price improvement, we sum over all partial executions as follows:

$$PrcImprove = q \times \sum_{i=1}^{n} \frac{VolTrade_i}{Volume} \times \frac{Price - Mid_t}{Mid_t}$$

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, Mid_t is the midpoint of the best bid and ask prices at the time of order entry, t, $Price$ is the execution price, $VolTrade_i$ is the volume executed for the ith partial fill, and $Volume$ is the number of shares entered in the original limit order.

If the price moves away to a more unfavorable level and the order fails to execute, the limit order trader can either amend the order or delete the order from the limit order book. 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:

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

where $VolFail$ is the number of shares that fail to trade and Mid_{t+r} is the midpoint of the best bid and ask prices at the time of order amendment or cancellation. We assume that all outstanding orders remaining in the limit order book at the end of the trading day are cancelled.

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

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

Finally, to arrive at daily measure of *LTC*, we weight each order by the order size for each trading day and trader type.[11]

Table 2, Panel A shows the summary statistics for total *LTC*. Limit order strategies are costly for institutional investors. The average limit order submitted by an institutional broker incurs a cost of 0.826 basis points while *HFT* and *Retail* receive benefits of 0.247 and 0.641 basis points, respectively, for their limit orders.

Table 2, Panels B and C decompose total *LTC* into the benefits of price improvement and the costs of non-execution, respectively. Comparing between the trader types, we find that retail investors receive the largest amount of price improvement for their limit orders. However, retail investors also face the highest non-execution costs. We find that institutional brokers suffer the

[11] Our findings remain qualitatively similar if we equally weight LTC across the orders.

highest non-execution costs at 2.83 basis points for the average limit orders submitted to either the best bid or ask price. In the next part of the analysis, we investigate how *LTC* changes for each of the broker categories when HFT are granted a larger speed advantage.

4.2 Results

We use a difference-in-differences estimator to formally examine the impact of trading speeds on HFT and non-HFT trading costs. The difference-in-differences approach offers several advantages over standard event study methodologies (Roberts and Whited, 2012). Specifically, this approach overcomes the problem of omitted trends while also controlling for the unobserved differences between two different groups of firms by analysing the same firms before and after the change (Roberts and Whited, 2012). In our application, the difference-in-differences methodology controls for market-wide trends in limit order transaction costs that affect both HFT and non-HFT.

The key assumption for the difference-in-differences estimator is the parallel trends assumption, which requires the same trends in outcomes for the treatment and control groups prior to treatment (Roberts and Whited, 2012). For our analysis, we assume that in the absence of ITCH, market wide factors influencing transaction costs for HFT and non-HFT behave in the same manner over the sample period. We conduct robustness tests in Section 5.0 to ensure that the parallel trends assumption holds.

Table 3 reports the results for the difference-in-differences regressions in which HFT represents the control group and non-HFT represents the treatment group. The dependent variable is limit order transaction costs (LTC). We also include stock fixed effects to control for unobserved firm effects. For the full stock sample, the main variable of interest, Non-HFT ×

Post-ITCH, is positive and significant indicating that non-HFT *LTC* increases, relative to HFT, after the implementation of the faster data speed. Specifically, we find that institutional and retail *LTC* increase by 1.73 and 58.6 basis points, respectively, once trading becomes faster for HFT.

Faster trading speeds could have a different impact on large and small stocks. Trading in highly liquid, larger stocks tend to be highly competitive and it may be difficult to exploit speed advantages in this subsample of stocks. On the other hand, small stocks have higher information asymmetry, meaning that fast traders could potentially gain more from their speed advantage in smaller, less liquid stocks. Table 3, columns 4 to 7 show the results for the large and small stock subsamples. For *Retail*, *LTC* increases for both large and small stock sub-samples. The magnitude of the change is much larger for the smaller stocks, reflecting the higher levels of information asymmetry in these stocks. For *Institutional*, we find that *LTC* increase by 4.2 basis points for small stocks while we detect no change in trading costs for large stocks. To gain a better understanding of why *LTC* changes pre- and post- the faster data feed, we decompose *LTC* into the individual components, *Non-execution* and *PrcImprove*.

Because total *LTC* captures both the benefits of price improvement and the costs of non-execution, we cannot determine whether costs change due to lower levels of price improvement or larger implementation shortfalls from aggregate measures of *LTC*. Table 4, Panel A, which reports the difference-in-difference results for non-execution costs, shows that the costs of implementation shortfalls rise for both *Institutional* and *Retail* after the adoption of ITCH. The result is consistent across the full stock sample as well as the large and small stock subsamples. Table 4, Panel B presents the results for *PrcImprove*. Because *PrcImprove* is stated in terms of costs, an increase in *PrcImprove* post-ITCH indicates that traders are receiving less price improvement for their limit orders. For retail orders, we find that *PrcImprove* worsens for both large and small stock subsamples. However, the results for *PrcImprove* are weaker than those of *Non-execution*. For institutional orders, our results are mixed. While the costs of non-execution increase for *Institutional* in both large and small stocks, we find that their limit orders receive more *PrcImprove* in large stocks, which explains the insignificant result for total *LTC* from Table 3.

Together our results show that the costs of non-execution are higher for institutional and retail brokers when HFT trading speeds improve, which increases the total costs of non-HFT limit order strategies. In Section 5.0, we investigate how HFT activity could affect the execution probabilities of non-HFT limit orders.

4.3 Robustness and additional tests

4.3.1 Strategic institutional trading

To conceal trading intentions, institutions often use algorithms to break up a single large order into multiple smaller orders (O'Hara, 2015). For example, many institutions rely on volume-weighted average price (VWAP) or time-weighted average price (TWAP) algorithms to minimise the trading costs of a much larger parent order. Alternatively, institutional brokers could have their own proprietary HFT strategies. For these reasons, small institutional order flow may be more difficult to detect regardless of the speed advantage, or alternatively may also contain HFT orders. Given that small institutional orders are more difficult to detect, we expect costs for large institutional orders to increase more, compared to the costs for smaller institutional orders. To investigate further, for each stock, we rank all institutional orders entered at the best bid or ask prices by size and further separate these orders into large institutional (top quintile) and small institutional (bottom quintile) orders.

Table 5 presents the results for the difference-in-difference regressions for limit order transaction costs based on institutional order size quintile subsamples. Consistent with our predictions, we find that *LTC* is higher after the implementation of ITCH for the largest institutional order quintile, as HFT algorithms can detect large orders more easily. For the smaller quintiles, we find a decrease in *LTC* indicating that institutional brokers are also becoming more strategic in their limit order placements. Alternatively, smaller institutional orders could also contain HFT order flow from institutional brokers or their clients. We investigate this possibility further in the next subsection.

4.3.2 Alternative HFT proxy

To avoid contamination by non-HFT order flow, our main analysis uses an HFT sub-sample that includes only orders submitted by pure proprietary HFT firms. However, it is possible that institutional brokers also implement proprietary HFT strategies and thus, some institutional order flow may be incorrectly classified as non-HFT. Because high frequency trading strategies are characterized by very high number of small orders to generate small profits per trade (Gomber et al., 2011), we use the smallest quintile of institutional orders as an alternative proxy for HFT. In Table 6, we compare small institutional LTC against larger institutional and retail LTC before and after the introduction of ITCH. The results are largely consistent with our main results. Specifically, we find that Non-HFT trading costs increase, relative to the trading costs for our alternative HFT proxy, after the decrease in latencies.

4.3.3 Falsification tests for alternative sample period

The introduction of ITCH on the ASX represents the only shock to trading speeds in recent years.[12] Thus, we should not observe a rise in non-HFT execution costs, relative to HFT costs, over a different sample period. To ensure that the parallel trends assumption holds for our analysis, we estimate a difference-in-difference regression for our sample stocks over the period January 1, 2011 to March 31, 2011 (pre-event) and May 1, 2011 to July 31, 2011 (post-event), which is precisely one year prior to our sample period. Table 7 shows that the interaction variable, *Non-HFT* × *Post-ITCH*, is insignificant for all model specifications. We fail to observe a similar rise in institutional or retail *LTC*, further supporting our finding that non-HFT limit order strategies become more expensive as a result of faster HFT.

4.3.4 Removing control variables

As a further test of the parallel trends assumption, we repeat the difference-in-differences regressions after removing the control variables. In unreported results, we find that the magnitudes and significance of the coefficients remain similar. These results indicate that HFT and non-HFT transaction costs respond to control variables in a similar manner before and after the implementation of ITCH.

5.0 HFT strategies

5.1 Predatory HFT proxy

Previous studies document a strong relationship between trade imbalances and future returns (Chordia, Roll and Subrahmanyam, 2002; Chordia and Subrahmanyam, 2004). Using

[12] The ASX introduced co-location in 2010. Because co-location was taken up by individual broker-dealers progressively over many months, we do not use co-location for our analysis.

more granular limit order book data, more recent studies also find strong evidence that order imbalances between the buy and sell schedules of the limit order book are significantly related to future stock returns (Cao, Hansch and Wang, 2008; Cont, Kukanov and Stoikov, 2013). Cont, Kukanov and Stoikov (2013) show that high-frequency price changes are mainly driven by imbalances between supply and demand at the best bid and ask prices. Specifically, large buying (selling) pressure on the bid (ask) price predicts future price rises (falls). Further, Ronaldo (2004) examines how the state of the limit order book can affect a trader's order submission strategy. We use the information contained in the state of the limit order book to proxy for predatory HFT trading.

Fig. 1 shows the scenario when there is a large depth imbalance on the bid side of the order book. The green boxes represent buy orders and the red box represents a sell order. In anticipation of a higher future price, a strategic limit order trader would want to be at the front of the limit order queue (Fig. 1a). In this case, the buy limit order will execute against the next incoming sell market order, and the broker will benefit from the predicted future price rise. If the trader does not have a buy order existing at the front of the limit order queue, the trader can either wait for their order to execute, with the risk that the price moves away, or trade strategically by submitting a market order and incurring the cost of the bid-ask spread (Fig. 1b). Conversely, because the predicted price is likely to be higher due to the demand-supply imbalance, a limit order on the ask side of the book faces high adverse selection risk and a strategic trader is likely to reduce this risk by cancelling the order (Fig. 1c). In comparison to slower traders, fast traders can submit market orders and cancel their limit orders more quickly, which decreases their adverse selection risk, at the slower trader's expense. Reducing trading latencies give fast traders an even greater advantage in monitoring their orders. We use the shape of the limit order book at the time of a trade or order cancellation to proxy for predatory HFT behaviour.

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 prices and the volume available at the best ask prices 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 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. To avoid capturing the volume of the trade itself, we compute *DepthImbalance* for the three bid and ask price levels behind the price of the trade. For robustness, we repeat the analysis for *DepthImbalance* calculated based on the volume available five price levels behind the trade. We arrive at a daily measure of *DepthImbalance* by volume weighting over the number of shares corresponding to each submitted order.

Table 8 presents summary statistics for *DepthImbalance* for HFT, institutional and retail broker types. For trade executions, a large value of *DepthImbalance* indicates that the trader is trading strategically based on the state of the limit order book. Conversely for cancellations, a small value of *DepthImbalance* indicates that the broker is successful in reducing their adverse selection risk.

Table 8, Panel A show that *HFT* are the most successful in their limit order strategies. Comparing across the trader types, *DepthImbalance* is highest for *HFT*, indicating that their limit orders are successfully executing when prices are supported by large depth in the favorable direction. Institutional limit orders suffer the highest adverse selection, indicated by the negative value for *DepthImbalance*. A negative *DepthImbalance* indicates that there is more depth in the opposite side of the order book, which is likely to reflect in an unfavorable price change.

Similarly, for market orders in Table 8, Panel B, we find that HFT are most strategic in their order placement strategies. Specifically, we find that HFT submit buy (sell) market orders when there is more depth on the bid (ask) prices, relative to the ask (bid) prices. Thus, HFT actively buy before a predicted price rise and sell prior if a negative return is forecasted.

Turning to cancellations in Table 8, Panel C, we also find that HFT are most strategic in their cancellation strategies. The lower values for *DepthImbalance* indicate that HFT cancel orders when there is only a small amount of depth supporting favorable future price movements. In other words, relative to other broker types, HFT are cancelling buy (sell) limit orders when there is a low amount of depth on the bid (sell) side of the order book to support a higher (lower) future price.

Our results also have implications for the market making role of HFT in equity markets. The previous literature generally concludes that HFT market making improves by reducing spreads and increasing depths (Hasbrouck and Saar, 2013). However, these findings are typically based on traditional measures of market depth, aggregated across both bid and ask prices (see Hasbrouck and Saar, 2013; Degryse, de Jong and van Kervel, 2011). Aggregated measures of market depth do not capture 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. Additionally, Van Kervel (2014) notes that 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.

In contrast to traditional market makers, HFT market makers have no obligation to provide liquidity on both sides of the order book, leading to concerns that HFT market making could result in market instability (see O'Hara, 2015). Our results support these concerns. Table 8, Panels A and B shows large *DepthImbalance* values for HFT, indicating that HFT systematically supply liquidity to the thick side of the order book, where it is least needed, and demand liquidity from the thin side of the order book, where it is most needed. Specifically, HFT limit buy orders execute at times when there is a large depth imbalance towards the bid side of the limit order book. Similarly, their market buy orders execute against sell limit orders when the ask side of the limit order book is already thin, relative to the bid side of the book.

5.2 Results

We conduct difference-in-difference regressions for *DepthImbalance* around the introduction of ITCH. Reducing data feed latencies allows fast traders to monitor and submit orders more efficiently. Thus, we expect *DepthImbalance* to improve for HFT limit and market order executions. Similarly, if HFT are using their speed advantage to cancel orders, we expect to observe a decrease in HFT *DepthImbalance* after the implementation of ITCH.

Table 9, Panel A presents the difference-in-difference results for limit orders. The positive coefficient on *HFT* indicates that HFT limit orders, compared to non-HFT limit orders, are on average executing when there is more depth supporting a favorable future price. Importantly, the positive coefficient on *HFT* × *Post-ITCH* indicates that HFT are even more successful in their limit order executions after their trading latencies decrease. The result is consistent across both small and large stock subsamples.

We find similar results tor market orders in Table 9, Panel B. Specifically, in both pre- and post-ITCH periods, HFT are more successful than non-HFT at demanding liquidity from the thin side of the limit order book. For the full sample of stocks, HFT become more predatory after trading speeds increase, indicated by the positive coefficient on the interaction term. However, faster speeds only provide an advantage to HFT in the large stock subsample. In contrast, for cancellations in Table 9, Panel C, we find that HFTs are more strategic in their limit order cancellations, but the result is only significant for small stocks.

Overall, these results show that predatory HFT strategies are increasing the non-execution costs of non-HFT limit orders. Faster data feeds enable HFTs to process information in the limit order book more efficiently than slower traders. Our results show that this speed advantages enables HFTs to trade ahead of non-HFT limit orders more frequently, when it is beneficial to do so.

6.0 Conclusion

We contribute to the debate on the benefits and concerns related to HFT by investigating the channels through which HFT activity affects the transaction costs of non-HFT investors. In 2012, the Australian Securities Exchange (ASX) implemented ITCH, an ultra-low latency protocol for accessing ASX market information, which was selectively taken up by brokers relying on low latency strategies. The introduction of ITCH on the ASX provides a clean natural experiment to test whether faster trading speeds benefiting one group of traders comes at a cost to another group of slower traders. We use difference-in-difference analysis around the implementation of ITCH to isolate the effects of faster trading speeds on the cost of HFT and non-HFT limit order strategies.

We find that faster trading speeds result in a wealth transfer from slow traders to fast traders. Specifically, we find that institutional and retail limit order trading costs increase, relative to the costs for HFT, after the reduction in trading latencies. Decomposing limit order costs into the benefits of price improvement and the costs of non-execution, we show that the increase in limit order trading costs is due to higher implementation shortfalls. Thus, the probability of a non-HFT limit order successfully executing falls when HFTs become even faster. When their limit orders fail to execute, non-HFTs are likely to resubmit their limit order at a less favorable price, which increases their overall execution costs.

Faster data feed speeds enable low latency traders to respond to changes in the state of the limit order book more efficiently. Using order book depth imbalance at the time of a trade or order cancelation as a proxy for predatory HFT activity, we find that HFTs become more strategic after the reduction in trading speeds. Specifically, HFTs are more successful in buying ahead of future price rises and selling ahead of future price falls. HFTs also become more successful in reducing their risk of adverse selection by selectively cancelling existing limit orders that are in the opposite direction to future price movements. Our results show that predatory trading is one mechanism through which speed benefits fast traders at the expense of slower traders.

Our results also have implications for market quality. In contrast to the market making role of HFT documented in the previous literature, we find that HFT typically supply liquidity to the thick side of the limit order book, where it is not required. In anticipation of future price movements, HFT strategically demand liquidity from the thin side, where liquidity is most needed, thus contributing to market instability.

The regulatory landscape for high frequency trading is rapidly changing, with regulators contemplating order cancellation fees, minimum order exposure times and transaction taxes to curb the growth of HFT. Our study provides one mechanism through which faster trading speeds could be detrimental to the quality of our financial markets.

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*.

Cao, C., Hansch, O. and Wang, X. (2009) 'The information content of an open limit-order book', *Journal of Futures Markets,* 29(1), 16-41.

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

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.

Cont, R., Kukanov, A. and Stoikov, S. (2014) 'The price impact of order book events', *Journal of Financial Econometrics,* 12(1), 47-88.

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.

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

Hagströmer, B. and Nordén, L. (2013) 'The diversity of high-frequency traders', *Journal of Financial Markets,* 16(4), 741-770.

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.

Hoffmann, P. (2014) 'A dynamic limit order market with fast and slow traders', *Journal of Financial Economics*, 113(1), 156-169.

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

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

O'Hara, M. (2015) 'High frequency market microstructure', *Journal of Financial Economics*, Forthcoming.

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

Roberts, M. R. and Whited, T. M. (2012), Endogeneity in Empirical Corporate Finance. Simon School Working Paper No. FR 11-29.

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

Tables and Figures

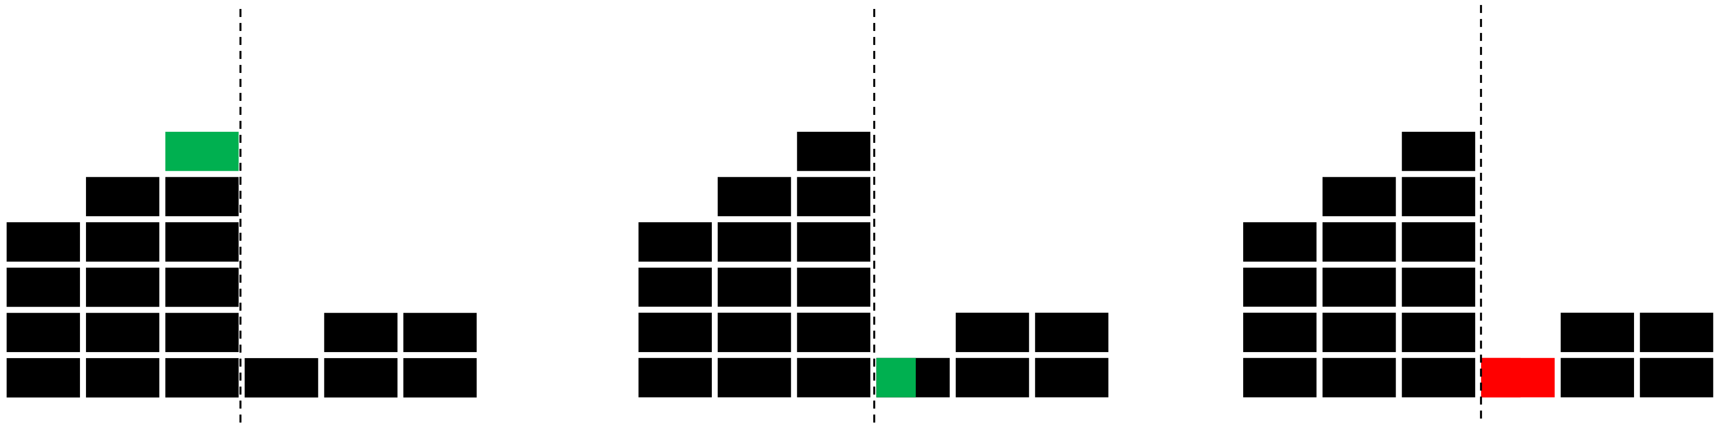

Fig. 1. Predatory HFT activity
Fig. 1 shows a limit order book where there is a large depth imbalance on the bid side, relative to the ask side of the limit order book. The dotted line represents the bid-ask midpoint. The green boxes in Figures 1a to 1b represent a buy limit and market order, respectively. The red box in Figure 1c represents a sell limit order.

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. We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. Panel B reports measures of average daily trading activity for HFT, institutional and retail brokers. All other variables are measured daily before averaging across the stocks in the sample. The final sample contains 94 stocks.

	Mean	Std. Dev.	Q1	Median	Q3
Panel A: Stock characteristics					
Market capitalization (bil.)	13.52	22.77	2.844	10.00	114.8
Price	11.02	12.61	2.930	5.670	14.21
Volume	4,013,013	5,245,714	1,116,000	2,720,000	4,137,000
Trades	2,189	1,619	1,165	1,629	2,389
Trade size	1,584	2,452	287.3	709.1	1,881
Volatility	0.022	0.007	0.018	0.020	0.024
QSpread (cents)	1.023	0.370	0.946	0.991	1.110
Panel B: Trader characteristics					
Volume traded					
HFT	123,016	150,245	38,470	82,340	134,300
Institutional	2,447,549	2,857,247	731,000	1,654,000	2,816,000
Retail	274,566	662,426	46,680	96,610	170,400
Trade size					
HFT	1,722	2,999	386.2	675.0	1,799
Institutional	1,544	2,280	276.1	637.0	1,903
Retail	2,554	3,989	441.8	1,058	2,615

Table 2.

Summary statistics – Limit order transaction costs

Table 2 reports statistics for limit order transaction costs (*LTC*). We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. *LTC* is calculated as:

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

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, Mid_t is the midpoint of the best bid and ask prices at time of order entry, t, Mid_{t+r} is the midpoint of the best bid and ask prices at the time of order amendment or cancellation, and $Price$ is the execution price. $VolTrade$ is the volume executed for the *i*th partial fill, $VolFail$ is the number of shares that fail to trade and $Volume$ is the number of shares entered in the original limit order. To arrive at a daily measure, we weight by order size for each trading day and trader type.

	Mean	Std. Dev.	Q1	Median	Q3
Panel A: Total LTC					
HFT	-0.247	0.921	-0.807	-0.188	0.288
Institutional	0.826	0.805	0.506	0.938	1.219
Retail	-0.641	5.431	-4.678	-0.493	0.546
Panel C: PrcImprove					
HFT	-2.644	2.071	-3.777	-2.449	-1.020
Institutional	-2.002	1.456	-2.619	-1.793	-1.029
Retail	-4.924	6.673	-8.901	-4.779	-2.449
Panel B: Non-execution					
HFT	2.397	1.741	1.211	2.187	3.267
Institutional	2.828	1.332	2.074	2.680	3.278
Retail	3.320	2.277	2.032	3.188	4.286

Table 3.

Difference-in-difference regressions for limit order transaction costs

Table 3 reports difference-in-difference regression results for limit order transaction costs (*LTC*). We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. The dependent variable, *LTC*, measures the daily volume-weighted limit order transaction costs for each trader type in each stock. *Non-HFT* is an indicator variable equal to 1 for the non-HFT category indicated in the column heading, and 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Price* is the daily volume weighted average price. *Total Volume* is the total daily volume of shares transacted by all traders in the stock. *Volatility* is the difference between the daily high and low prices, divided by the average time-weighted midpoint price. *QSpread* is the daily time-weighted quoted spread. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median on January 1, 2012. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All		Large stocks		Small stocks	
	Institutional	Retail	Institutional	Retail	Institutional	Retail
Non-HFT	-2.259 ***	-129.761 ***	-0.618 **	-56.171 ***	-4.295 ***	-253.882 ***
	(0.402)	(9.463)	(0.314)	(7.838)	(0.975)	(27.654)
Post - ITCH	2.481	-18.555	-0.173	-13.024	4.998	25.051
	(2.315)	(55.652)	(1.907)	(48.306)	(5.142)	(148.107)
Non-HFT × Post-ITCH	1.726 ***	58.646 ***	0.533	23.825 **	4.232 ***	127.755 ***
	(0.529)	(12.460)	(0.432)	(10.784)	(1.226)	(34.903)
Log(Price)	0.236 **	4.167	0.007	1.091	1.675 ***	52.739 ***
	(0.118)	(2.752)	(0.072)	(1.786)	(0.645)	(18.371)
Log(Total Volume)	-1.544 ***	-1.965	-0.83 ***	-1.011	-2.596 ***	4.566
	(0.291)	(6.936)	(0.304)	(7.588)	(0.571)	(16.626)
Volatility	1.335 ***	7.518 **	0.369 ***	1.973	2.086 ***	15.233 **
	(0.135)	(3.167)	(0.126)	(3.147)	(0.258)	(7.248)
QSpread	0.027 *	1.039 ***	0.05 ***	0.46	0.032	2.02 **
	(0.015)	(0.343)	(0.013)	(0.318)	(0.029)	(0.826)
Constant	13.81 ***	-60.44	8.978 *	-31.74	20.19 **	-443.8
	(4.823)	(115.300)	(5.154)	(129.100)	(9.391)	(273.800)
Obs.	16,868	16,250	6,472	6,432	4,858	4,540
Adj. R-square	0.0359	0.0411	0.0233	0.0192	0.0563	0.0557

Table 4.

Difference-in-difference regressions for the components of limit order transaction costs

Table 4 reports difference-in-difference regression results for the two components of limit order transaction costs (*LTC*). We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. In Panel A, the dependent variable, *Non-execution costs*, measures the daily volume-weighted cost of non-executed limit orders for each trader type in each stock. In Panel B, the dependent variable, *PrcImprove*, measures the daily volume-weighted gains from executed limit orders for each trader type in each stock. *PrcImprove* is expressed as a cost and thus, a negative value of *PrcImprove* is interpreted as a gain to the limit order trader. *Non-HFT* is an indicator variable equal to 1 for the non-HFT category indicated in the column heading, and 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Price* is the daily volume weighted average price. *Total Volume* is the total daily volume of shares transacted by all traders in the stock. *Volatility* is the difference between the daily high and low prices, divided by the average time-weighted midpoint price. *QSpread* is the daily time-weighted quoted spread. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median on January 1, 2012. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All		Large stocks		Small stocks	
	Institutional	Retail	Institutional	Retail	Institutional	Retail
Panel A: Non-execution costs						
Non-HFT	-2.269 ***	-42.166 ***	-1.714 ***	-10.673 ***	-2.013 ***	-104.736 ***
	(0.229)	(4.062)	(0.205)	(1.565)	(0.402)	(13.601)
Post - ITCH	1.736	7.832	-1.951	-9.553	6.475 ***	50.325
	(1.321)	(23.887)	(1.248)	(9.643)	(2.118)	(72.840)
Non-HFT × Post-ITCH	1.902 ***	27.32 ***	1.547 ***	8.243 ***	1.96 ***	69.139 ***
	(0.302)	(5.348)	(0.282)	(2.153)	(0.505)	(17.166)
Log(Price)	0.09	1.419	-0.038	0.012	0.299	21.678 **
	(0.068)	(1.181)	(0.047)	(0.356)	(0.266)	(9.035)
Log(Total Volume)	-0.163	-2.925	-0.100	1.694	-0.265	-11.295
	(0.166)	(2.977)	(0.199)	(1.515)	(0.235)	(8.177)
Volatility	1.063 ***	5.473 ***	0.336 ***	-0.107	1.483 ***	12.031 ***
	(0.077)	(1.359)	(0.083)	(0.628)	(0.106)	(3.565)
QSpread	0.007	0.317 **	0.035 ***	0.062	-0.007	0.712 *
	(0.008)	(0.147)	(0.008)	(0.064)	(0.012)	(0.406)
Constant	1.042	0.3865	2.085	-21.61	-2.954	-26.87
	(2.752)	(49.510)	(3.371)	(25.760)	(3.869)	(134.700)
Obs.	16,868	16,250	6,472	6,432	4,858	4,540
Adj. R-square	0.041	0.032	0.029	0.032	0.078	0.042

Table 4—Continued

Panel B: PrcImprove						
Non-HFT	0.341	-66.561 ***	1.207 ***	-32.343 ***	-1.637 **	-113.302 ***
	(0.223)	(5.960)	(0.151)	(4.857)	(0.736)	(15.189)
Post - ITCH	0.657	-16.758	1.809 **	-5.252	-1.197	10.606
	(1.286)	(35.055)	(0.915)	(29.935)	(3.882)	(81.345)
Non-HFT × Post-ITCH	-0.336	23.947 ***	-1.113 ***	12.221 *	1.749 *	47.38 **
	(0.294)	(7.848)	(0.207)	(6.683)	(0.925)	(19.170)
Log(Price)	0.127 *	1.923	0.039	0.457	1.158 **	22.936 **
	(0.066)	(1.734)	(0.034)	(1.107)	(0.487)	(10.090)
Log(Total Volume)	-1.17 ***	1.910	-0.697 ***	-1.101	-1.962 ***	14.737
	(0.162)	(4.369)	(0.146)	(4.702)	(0.431)	(9.131)
Volatility	0.161 **	0.325	0.027	1.357	0.423 **	1.029
	(0.075)	(1.995)	(0.061)	(1.950)	(0.195)	(3.981)
QSpread	0.015 *	0.545 **	0.011 *	0.431 **	0.027	0.76 *
	(0.008)	(0.216)	(0.006)	(0.197)	(0.022)	(0.454)
Constant	10.4 ***	-73.6	6.685 ***	-29.19	19.57 ***	-372.3 **
	(2.679)	(72.650)	(2.471)	(79.980)	(7.090)	(150.400)
Obs.	16,868	16,250	6,472	6,432	4,858	4,540
Adj. R-square	0.081	0.029	0.103	0.017	0.072	0.041

Table 5.

Difference-in-difference regressions for limit order transaction costs – Institutional order size quintiles

Table 5 reports difference-in-difference regression results for limit order transaction costs (*LTC*) for institutional orders. We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. We further sort institutional orders into quintiles based on the size of the submitted order. The dependent variable, *LTC*, measures the daily volume-weighted limit order transaction costs for each institutional order size quintile in each stock. *Non-HFT* is an indicator variable equal to 1 for the institutional order size category indicated in the column heading, and 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Price* is the daily volume weighted average price. *Total Volume* is the total daily volume of shares transacted by all traders in the stock. *Volatility* is the difference between the daily high and low prices, divided by the average time-weighted midpoint price. *QSpread* is the daily time-weighted quoted spread. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median on January 1, 2012. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	1 (smallest)	2	3	4	5 (largest)
Non-HFT	0.378 ***	0.269 **	0.348 ***	0.389 ***	-2.142 ***
	(0.105)	(0.109)	(0.108)	(0.106)	(0.418)
Post - ITCH	0.656	0.608	0.694	0.518	2.042
	(0.608)	(0.625)	(0.621)	(0.612)	(2.406)
Non-HFT × Post-ITCH	-0.52 ***	-0.325 **	-0.291 **	-0.085	1.609 ***
	(0.139)	(0.143)	(0.142)	(0.140)	(0.549)
Log(Price)	0.079 **	0.09 ***	0.09 ***	0.107 ***	0.223 *
	(0.031)	(0.032)	(0.032)	(0.031)	(0.123)
Log(Total Volume)	-0.221 ***	-0.269 ***	-0.223 ***	-0.363 ***	-1.702 ***
	(0.076)	(0.079)	(0.078)	(0.077)	(0.303)
Volatility	0.555 ***	0.617 ***	0.611 ***	0.676 ***	1.332 ***
	(0.036)	(0.037)	(0.036)	(0.036)	(0.141)
QSpread	-0.004	-0.004	-0.004	-0.005	0.034 **
	(0.004)	(0.004)	(0.004)	(0.004)	(0.015)
Constant	0.547	1.291	0.469	2.436 *	16.01 ***
	(1.266)	(1.303)	(1.294)	(1.276)	(5.011)
Obs.	16,866	16,866	16,866	16,868	16,868
Adj. R-square	0.044	0.049	0.054	0.058	0.035

Table 6.

Difference-in-difference regressions for limit order transaction costs – Small institutional orders as alternative HFT proxy

Table 6 reports difference-in-difference regression results for limit order transaction costs (*LTC*) using small institutional orders to proxy for HFT. We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. In this table, we further sort institutional orders into quintiles based on the size of the submitted order. The dependent variable, *LTC*, measures the daily volume-weighted limit order transaction costs for each trader type in each stock. *Non-HFT* is an indicator variable equal to 1 for the non-HFT category indicated in the column heading, and 0 for institutional orders belonging in size quintile 1 (smallest orders). *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Price* is the daily volume weighted average price. *Total Volume* is the total daily volume of shares transacted by all traders in the stock. *Volatility* is the difference between the daily high and low prices, divided by the average time-weighted midpoint price. *QSpread* is the daily time-weighted quoted spread. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median on January 1, 2012. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All		Large stocks		Small stocks	
	Institutional (Quintiles 2-5)	Retail	Institutional (Quintiles 2-5)	Retail	Institutional (Quintiles 2-5)	Retail
Non-HFT	-0.63 ***	-32.19 ***	-0.432 ***	-7.258 ***	-0.42 **	-64.213 ***
	(0.123)	(3.111)	(0.149)	(1.324)	(0.186)	(8.800)
Post - ITCH	0.997	2.065	-0.270	-4.951	1.423	17.086
	(0.698)	(18.018)	(0.842)	(7.606)	(1.046)	(50.996)
Non-HFT × Post-ITCH	0.992 ***	19.671 ***	0.808 ***	5.725 ***	1.205 ***	34.825 ***
	(0.173)	(4.349)	(0.208)	(1.856)	(0.260)	(12.283)
Log(Price)	0.008	0.901	-0.032	-0.045	0.001	7.79 *
	(0.035)	(0.879)	(0.033)	(0.291)	(0.087)	(4.107)
Log(Volume)	0.039	0.382	0.244	1.973	0.204	-2.706
	(0.096)	(2.442)	(0.149)	(1.326)	(0.127)	(6.126)
Volatility	0.702 ***	4.13 ***	0.236 ***	-(0.2110)	0.909 ***	8.769 ***
	(0.046)	(1.147)	(0.063)	(0.563)	(0.059)	(2.714)
QSpread	0.007 **	0.095	0.028 ***	0.072	0.005	0.180
	(0.003)	(0.086)	(0.006)	(0.053)	(0.004)	(0.209)
Constant						
Obs.	23,432	22,368	7,642	7,584	7,678	7,088
Adj. R-square	0.050	0.026	0.031	0.029	0.086	0.028

Table 7.

Difference-in-difference regressions for limit order transaction costs – Placebo tests

Table 5 reports difference-in-difference regression results for limit order transaction costs (*LTC*). We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2011 to March 31, 2011 (*Pre-ITCH*) and May 1, 2011 to July 31, 2011 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. The dependent variable, *LTC*, measures the daily volume-weighted limit order transaction costs for each trader type in each stock. *Non-HFT* is an indicator variable equal to 1 for the non-HFT category indicated in the column heading, and 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Price* is the daily volume weighted average price. *Total Volume* is the total daily volume of shares transacted by all traders in the stock. *Volatility* is the difference between the daily high and low prices, divided by the average time-weighted midpoint price. *QSpread* is the daily time-weighted quoted spread. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All		Large stocks		Small stocks	
	Institutional	Retail	Institutional	Retail	Institutional	Retail
Non-HFT	3.029 ***	-76.512 ***	1.878 ***	-37.529 ***	7.681 ***	-161.835 ***
	(0.458)	(9.252)	(0.414)	(7.919)	(1.291)	(33.902)
Post - ITCH	4.913	63.416	3.825	131.94 **	20.184 **	-80.141
	(3.396)	(70.040)	(2.942)	(56.177)	(8.931)	(228.977)
Non-HFT × Post-ITCH	-0.341	12.15	-0.559	-13.833	-2.002	78.804
	(0.669)	(13.517)	(0.590)	(11.269)	(1.977)	(52.283)
Log(Price)	0.145	2.992	0.121	-0.654	1.076	24.532
	(0.150)	(3.017)	(0.130)	(2.484)	(1.180)	(31.915)
Log(Volume)	0.081	3.800	0.859 **	25.782 ***	-1.982 *	-50.26 *
	(0.418)	(8.498)	(0.427)	(8.166)	(1.115)	(30.242)
Volatility	0.206 *	(1.7050)	(0.0620)	(0.0590)	1.288 **	(18.2150)
	(0.122)	(2.442)	(0.095)	(1.812)	(0.527)	(13.902)
QSpread	-0.002	-0.223	0.001	0.71 **	0.042	-1.182
	(0.015)	(0.303)	(0.017)	(0.332)	(0.043)	(1.132)
Constant						
Obs.	7,602	7,470	4,740	4,730	1,312	1,232
Adj. R-square	0.122	0.038	0.139	0.049	0.143	0.010

Table 8.

Summary statistics – Depth imbalance

Table 8 reports statistics for *DepthImbalance*, our proxy for HFT strategies. We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. *DepthImbalance* is calculated 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 three levels behind the best bid (ask) price immediately before order submission (Panels A and B) or cancellation (Panel C) 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 average over the orders submitted or cancelled for each stock and trader category.

	Mean	Std. Dev.	Q1	Median	Q3
Panel A: Limit orders					
HFT	0.106	0.078	0.039	0.109	0.156
Institutional	-0.004	0.039	-0.032	-0.015	0.012
Retail	0.000	0.012	-0.006	0.000	0.008
Panel B: Market orders					
HFT	0.031	0.041	0.006	0.026	0.053
Institutional	-0.009	0.041	-0.036	-0.003	0.022
Retail	-0.023	0.036	-0.036	-0.013	0.000
Panel C: Cancellations					
HFT	0.008	0.012	0.000	0.005	0.015
Institutional	0.030	0.015	0.020	0.029	0.038
Retail	0.012	0.015	0.003	0.010	0.022

Table 9.

Difference-in-difference regressions for HFT strategies

Table 9 reports difference-in-difference regression results for HFT strategies. We analyze order-level data for stocks in the ASX 100 for the periods January 1, 2012 to March 31, 2012 (*Pre-ITCH*) and May 1, 2012 to July 31, 2012 (*Post-ITCH*). Orders are classified into 4 categories based on the broker submitting the order: high-frequency trader (HFT), institutional, retail, and those not previously classified, other. For Panels A and B, we consider all top of book limit and market orders submitted at the best bid or ask price, respectively. For Panel C, we consider all limit orders priced at the best bid or ask price at the time of cancellation, regardless of whether the original order was submitted to the top of the book. We calculate *DepthImbalance* 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 three levels behind the best bid (ask) price immediately before order submission (Panels A and B) or cancellation (Panel C) 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 average over the orders submitted or cancelled for each stock and trader category. *Non-HFT* is an indicator variable equal to 1 for the non-HFT category indicated in the column heading, and 0 for HFT. *Post-ITCH* is an indicator variable equal to 1 if the observation occurs in the post-ITCH period, and 0 for the pre-ITCH period. *Large stocks* (*Small stocks*) refer to stocks with a market capitalization above (below) the median on January 1, 2012. All regressions control for stock fixed effects. We report heteroskedasticity-robust standard errors clustered by stock in parentheses. ***, ** and * indicate significance levels of 1%, 5% and 10%, respectively.

	All		Large stocks		Small stocks	
	Institutional	Retail	Institutional	Retail	Institutional	Retail
Panel A: Limit orders						
HFT	0.0445 ***	0.0578 ***	0.0413 ***	0.0507 ***	0.0437 ***	0.0743 ***
	(0.002)	(0.003)	(0.002)	(0.003)	(0.005)	(0.008)
Post - ITCH	0.0079 ***	-0.0086 ***	0.0083 ***	-0.0105 ***	0.009 *	0.010
	(0.002)	(0.003)	(0.002)	(0.003)	(0.005)	(0.007)
HFT × Post-ITCH	0.0195 ***	0.0361 ***	0.0194 ***	0.0379 ***	0.0316 ***	0.0313 ***
	(0.003)	(0.004)	(0.003)	(0.004)	(0.007)	(0.010)
Constant	-0.0077 ***	-0.0212 ***	-0.0097 ***	-0.0188 ***	-0.002	-0.0327 ***
	(0.001)	(0.002)	(0.002)	(0.002)	(0.004)	(0.005)
Obs.	17,810	17,320	7,490	7,392	4,146	3,946
Adj. R-square	0.101	0.099	0.150	0.137	0.096	0.103

Table 9—Continued

Panel B: Market orders						
HFT	0.0957 ***	0.0851 ***	0.0575 ***	0.0544 ***	0.1386 ***	0.1205 ***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.005)
Post - ITCH	-0.0101 ***	-0.0074 ***	-0.0116 ***	-0.0043 **	-0.0096 ***	-0.0086 *
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.005)
HFT × Post-ITCH	0.0106 ***	0.0082 ***	0.0231 ***	0.0158 ***	0.0034	0.003
	(0.002)	(0.003)	(0.003)	(0.003)	(0.005)	(0.006)
Constant	-0.0064 ***	0.0043 ***	-0.0022 *	0.001	-0.0125 ***	0.0065 **
	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)
Obs.	19,892	19,772	7,512	7,508	5,572	5,486
Adj. R-square	0.284	0.184	0.292	0.208	0.346	0.212
Panel C: Cancellations						
HFT	-0.005 *	-0.003	-0.0091 **	-0.004	-0.006	-0.001
	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.008)
Post - ITCH	-0.0048 *	0.0097 **	-0.003	0.006	0.000	0.0165 **
	(0.003)	(0.004)	(0.004)	(0.006)	(0.005)	(0.007)
HFT × Post-ITCH	-0.0012	-0.0131 **	0.0074	-0.0008	-0.0151 **	-0.0317 ***
	(0.004)	(0.006)	(0.006)	(0.008)	(0.007)	(0.011)
Constant	0.0163 ***	0.0136 ***	0.0123 ***	0.006	0.0169 ***	0.0149 ***
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)	(0.006)
Obs.	12,722	9,606	3,674	3,172	4,346	3,074
Adj. R-square	0.001	0.002	0.000	0.000	0.005	0.007