Equity Market Structure Literature Review

Part II: High Frequency Trading

By

Staff of the Division of Trading and Markets[1]

U.S. Securities and Exchange Commission

March 18, 2014

[1] This review was prepared by the Staff of the U.S. Securities and Exchange Commission. The Commission has expressed no view regarding the analysis, findings, or conclusions contained in this paper.

The following is Part II of a series reviewing recent economic literature on equity market structure. This SEC staff review summarizes those economic papers that analyze recent financial market data (2007 and later) and reach findings that in the view of staff of the Division of Trading and Markets are most relevant to important market structure issues facing the SEC.[2] Part I focused on papers that address market fragmentation – both visible and dark.[3] It also briefly noted the SEC's comprehensive review of equity market structure and gave an overview of the objectives of the staff's literature review.

Part II summarizes and discusses papers that address high frequency trading ("HFT"). These papers analyze non-public datasets in which market activity can be attributed to trading accounts that have been identified as engaging in HFT ("HFT Datasets"). A forthcoming Part III of the literature review will address a series of papers that do not have access to datasets in which market activity could be attributed to HFT accounts, but rather use various measures calculable from publicly available market data to proxy for HFT. Such HFT proxies include high message rates, bursts of order cancellations and modifications, high order-to-trade ratios, small trade sizes, and increases in trading speed. These proxies generally are associated with the broader phenomena of algorithmic trading and computer-assisted trading in all their forms.

As discussed in Section I.A below, the staff believes that focusing first on papers with datasets that specifically identify HFT account activity will be useful when subsequently assessing papers that use proxies for HFT, particularly the extent to which the various proxies in fact capture HFT activity as contrasted with other forms of algorithmic trading and computer-assisted trading. Part III also will address market structure performance in general, including volatility and investor transaction costs.

The staff's hope is that this literature review will help promote a dynamic exchange with and among the public, including investors, academics, securities industry participants, and others on the topic of HFT.

[2] This literature review does not include purely theoretical papers and also does not focus on the theoretical explanations of results often set forth in data papers. These theoretical papers and explanations are of great interest, and some are seminal papers that set up the economic foundation of many of the empirical papers. The primary objective of this literature review, however, is to set forth empirical results as a step in the staff's continued consideration of equity market structure issues.

[3] Part I is available on the SEC's equity market structure website at http://www.sec.gov/marketstructure/research.

Part Two:
High Frequency Trading

Table of Contents

I. HFT – Overview

A. Metrics for Defining HFT

The SEC's Concept Release on Equity Market Structure recognized that HFT is one of the most significant market structure developments in recent years.[4] It noted, for example, that estimates of HFT typically exceeded 50% of total volume in U.S.-listed equities and concluded that, "[b]y any measure, HFT is a dominant component of the current market structure and likely to affect nearly all aspects of its performance."[5]

The Concept Release also noted that the term "HFT" was not clearly defined. To deal with this problem, the Concept Release first generally defined "proprietary firm" as "professional traders acting in a proprietary capacity that generate a large number of trades on a daily basis."[6] These traders could be organized in a variety of ways, including as a proprietary trading firm (which may or may not be a registered broker-dealer and a member of FINRA), as the proprietary trading desk of a multi-service broker-dealer, or as a hedge fund.

Next, the Concept Release identified five other characteristics that often are attributed to HFT:

1. Use of extraordinarily high speed and sophisticated programs for generating, routing, and executing orders.

2. Use of co-location services and individual data feeds offered by exchanges and others to minimize network and other latencies.

3. Very short time-frames for establishing and liquidating positions.

4. Submission of numerous orders that are cancelled shortly after submission.

5. Ending the trading day in as close to a flat position as possible (that is, not carrying significant, unhedged positions overnight).

The Concept Release did not, however, suggest that <u>all</u> of these characteristics must be present for a proprietary firm to be properly classified as HFT. Such an approach may inappropriately narrow the range of firms that are classified as HFT. In this regard, Section II.A below notes that many of the HFT Dataset papers offer empirical findings that shed light on appropriate metrics for classifying particular trading accounts as HFT.

---

4 Securities Exchange Act Release No. 34-61358, 75 FR 3594, 3606 (January 21, 2010) ("Concept Release").

5 <u>Id</u>.

6 <u>Id</u>.

Moreover, in the absence of trading account data, the use of general proxies for HFT that can be calculated with publicly available, market-wide data may capture a great deal of algorithmic and computer-assisted trading that should not be classified as HFT. Examples of such HFT proxies derived from market-wide data include high message rates, bursts of order cancellations and modifications, high order-to-trade ratios, small trade sizes, and increases in trading speed. These market-wide proxies are associated with the broader phenomena of algorithmic trading and computer-assisted trading in all their forms. As discussed below, HFT represents a large subset, but by no means all, of algorithmic and computer-assisted trading.

For example, algorithmic trading encompasses a broad range of activity, including particularly the large order execution algorithms often used by or on behalf of institutional investors.[7] This type of algorithm takes institutional investor orders, which typically are too large to be executed all at once without excessive price impact, and slices them into many small orders that are fed into the marketplace over time. The staff notes that these large order algorithms should not be classified as HFT because they typically enable institutional investors to establish or liquidate positions with time horizons far beyond the primarily intraday horizons characteristic of HFT.

In addition, other types of computer-assisted trading tools are common in today's markets that may generate market activity that is difficult to distinguish from HFT, at least in the absence of datasets that can tie market activity to particular trading accounts. These tools include smart order routing systems that are designed to deal with the large number of trading venues in the fragmented U.S. equity market structure. They also include trading systems with automated functionalities that, while perhaps not falling within the definition of an algorithm (and therefore not appropriately classified as HFT), nevertheless enable orders to be submitted to the marketplace in ways that are far beyond the manual capacities of a human trader.

In sum, while HFT clearly is a large subset of algorithmic trading and computer-assisted trading, the HFT Dataset papers, as well as some recent market events and enforcement proceedings described below, indicate that other types of automated trading are significant and may be quite difficult to distinguish from HFT in the absence of trading account data that can be used to distinguish different types of market participants. Here are some examples:

- ASIC (2013) uses a dataset with trader identifiers to determine that 76% of small and fleeting orders, which had sparked concern in the Australian markets, were not in fact generated by HFTs. For example, the authors determined that almost 38% of the small and fleeting orders were generated by two non-HFT market participants engaging in buy-side strategies. They also estimated that at least 99.6% of all trading messages were sourced from an automated order processing program, indicating the

7 See Concept Release, 75 FR at 3602.

pervasive use of computer-assisted trading tools across all types of market participants, including those not classified as HFT. The paper further finds that firms classified as non-HFTs are almost as fast as HFTs in the Australian markets, and that non-HFTs have a higher tendency to amend orders within a short time period than HFTs.

- Hagstromer, Norden and Zhang (2013) examine the use of order types by different groups of traders on the NASDAQ OMS Stockholm exchange ("OMSX"), finding only small economic differences in order cancellation rates between HFTs and non-HFTs.

- Brogaard, Hagstromer et al. (2013) examine trading on OMXS with a dataset that includes member identifiers, as well as whether the members purchased varying levels of co-location services offering faster trading speeds. They find that numerous non-HFTs (according to two different HFT definitions) use co-location accounts, so use of these accounts does not necessarily mean that a market participant is an HFT.

- Clark-Joseph (2013) uses a dataset with trader identifiers for the E-Mini S&P 500 futures contract ("E-Mini") to classify 30 trading accounts as HFT. These 30 accounts included 8 large accounts that engaged in primarily aggressive strategies, in contrast to the primarily passive strategies that often generate a large number of order book messages. The 30 accounts classified as HFT represented 46.7% of total trading volume in the dataset, but initiated only 31.9% of all order entry, order modification, and order cancellation messages. For this dataset, an HFT proxy derived solely from order book messages would not capture a large segment of aggressive HFT.

- Brogaard et al. (2012) use a regulatory dataset with trader identifiers to study four technology enhancements on the London Stock Exchange that increased its speed of trading. In two of the four cases (including the one with the largest speed increase – six milliseconds), there was no significant impact on the HFT share of trading. For this dataset, an HFT proxy derived solely from an increase in trading system speed would not in fact reflect an increase in HFT in two of the four technology enhancements.

- *In re Knight Capital Americas LLC*, SEC Administrative Proceeding File No. 3-15570 (October 16, 2013) involved a significant error in an automated order routing system of a broker-dealer. Due to the error, the firm erroneously sent millions of orders into the market when it processed 212 small retail orders, generating more than 4 million trades in 45 minutes. In sum, the error did not arise from a trading strategy that could be classified as HFT.

- *In the matter of Hold Bros. Online Investment Services, LLC, et al.*, SEC Administrative Proceeding File No. 3-15046 (September 25, 2012) involved a broker-dealer that controlled many overseas traders who engaged in market manipulation. In one example of "layering," a single trader placed and cancelled 11 orders with different prices and sizes in less than one second. See also FINRA Letter of Acceptance, Waiver, and Consent No. 2010023771001, Re: Hold Brothers On-Line Investment Services, LLC. (September 25, 2012) (firm controlled 217 different trader groups with 2,432 identified traders that generated 400,000 trades per day during January through December 2009). Although it did not employ trading algorithms (and therefore was not HFT), the violative conduct did employ computer-assisted tools that enabled a rate of order entry and cancellations beyond what could be manually implemented and therefore could be mistaken for HFT activity.

These examples illustrate that all low latency and high frequency activity is not necessarily associated with proprietary firms that engage in what would appropriately be classified as HFT. Consequently, the staff believes that the HFT Dataset papers, with their focus on market activity that can be tied to identified HFT accounts, should be addressed in this Part II of the literature review. The broader phenomenon of algorithmic and computer-assisted trading will be addressed in a subsequent part of the literature review.

B. HFT Strategies

The Concept Release noted that the lack of a clear definition of HFT complicated the SEC's review of market structure. Proprietary firms may engage in a variety of different strategies, some of which may benefit market quality and some of which may be harmful. Rather than "attempt any single, precise definition of HFT,"[8] the Concept Release focused on the particular strategies and tools that may be used by proprietary firms and inquired whether any of them raised concerns that needed to be addressed.

In particular, the SEC requested comment on four types of short-term trading strategies – passive market making, arbitrage, structural, and directional.

Passive market making primarily involves the submission of non-marketable resting orders that provide liquidity to the marketplace at specified prices. The Concept Release noted that profits for these strategies do not depend primarily on directional price moves, but rather on earning a spread between bids and offers, as often supplemented by liquidity rebates offered by most markets for resting liquidity. Because these passive orders generally are not executed immediately and must rest on an order book, their prices may need to be updated frequently to reflect changing market conditions. The Concept Release noted that a passive market making strategy may generate an enormous number of order cancellations or modifications as orders are updated.

[8] 75 FR at 3607.

An arbitrage strategy generally seeks to capture pricing disparities between related products or markets, such as between an exchange traded product ("ETP") and its underlying basket of stocks. Arbitrage strategies also do not depend on directional price moves, but rather on price convergence.

Structural strategies attempt to exploit structural vulnerabilities in the market or in certain market participants. For example, traders with access to the lowest latency market data and trading tools may be able to profit by trading with market participants on a trading venue that is offering executions at stale prices.

Directional strategies generally involve establishing a long or short position in anticipation of a price move up or down. The Concept Release requested comment on two types of directional strategies – order anticipation and momentum ignition – that "may pose particular problems for long-term investors" and "may present serious problems in today's market structure."[9] An order anticipation strategy seeks to ascertain the existence of large buyers or sellers in the marketplace and then trade ahead of those buyers or sellers in anticipation that their large orders will move market prices (up for large buyers and down for large sellers). A momentum ignition strategy involves initiating a series of orders and trades in an attempt to ignite a rapid price move up or down. As noted in the Concept Release,[10] any market participant that manipulates the market has engaged in conduct that already is illegal. The Concept Release focused on the issue of whether additional regulatory tools were needed to address illegal practices, as well as any other practices associated with momentum ignition strategies.

The Concept Release noted that many of the strategies it discussed were not new, but that technology may allow proprietary firms to better identify and execute the strategies.[11] For example, it asked whether the speed of trading and the ability to generate a large number of orders might render momentum ignition strategies more of a problem today than in the past.[12] Comment generally was sought on the extent to which such strategies were used and their effect on market quality, particularly from the standpoint of long-term investors.

As discussed below, researchers of HFT have sought to address many Concept Release questions, as well as others.

II. Summary of Papers

A summary of the economic literature must begin by highlighting a formidable challenge facing any researcher of HFT – obtaining useful data that can identify HFT activity. Publicly available data on orders and trades does not reveal the identity of buyers and

---

9 75 FR at 3607, 3609.
10 75 FR at 3609.
11 75 FR at 3607.
12 75 FR at 3610.

sellers. As a result, it is impossible to identify orders and trades as originating from an HFT account when relying solely on publicly available information.

This review summarizes 31 papers for which the authors have managed to obtain access to datasets with non-public information that identifies, to a greater or lesser extent, activity arising from HFT accounts. These HFT Datasets can be divided into four categories: (1) data for equity trading on NASDAQ that NASDAQ has made available to researchers ("NASDAQ Datasets"); (2) data on trading in the E-Mini that the CFTC has made available to researchers ("E-Mini Datasets"); (3) data that was used by CFTC and SEC staff to prepare their report on the severe market disruption that occurred on May 6, 2010 ("Flash Crash") ("Flash Crash Datasets"); and (4) a variety of datasets made available to researchers by exchanges and regulators internationally ("International Datasets"). The particulars of the four categories of HFT Datasets are discussed in Section III.A below.

Perhaps the most noteworthy finding of the HFT Dataset papers is that HFT is not a monolithic phenomenon, but rather encompasses a diverse range of trading strategies. In particular, HFT is not solely, or even primarily, characterized by passive market making strategies that employ liquidity providing orders that rest on order books and can be accessed by others. For example, Carrion (2013) and Brogaard, Hendershott and Riordan (2013) analyze NASDAQ Datasets and find that more than 50% of HFT activity is attributable to aggressive, liquidity taking orders that trade immediately against passive resting orders. Moreover, the level and nature of HFT activity can vary greatly across different types of stocks. Brogaard, Hendershott and Riordan (2013) also find that HFTs are much less active in small-capitalization stocks than in large-capitalization stocks, but that 69% of their activity in small-capitalization stocks is attributable to aggressive orders.

The diversity of HFT highlights the importance of exercising care when using metrics to define HFT and proxies to associate with HFT. In particular, metrics and proxies that are based directly or indirectly on passive order book activity (such as message rates or cancellation rates) may have the effect of excluding a large volume of aggressive HFT activity. In addition, narrow metrics for other aspects of HFT activity, such as intraday and end-of-day inventory levels, can exclude a large volume of HFT activity. As noted in Section III.A.2 below, for example, the analysis by Kirilenko et al. (2011) of HFT activity in the E-Mini during the Flash Crash likely excluded at least 1/3rd of HFT activity.

Recognizing the diversity of HFT strategies also is essential when assessing the effect of HFT on market quality. As discussed in Section III.B below, different strategies can have quite varying effects on market quality. In general, the HFT Dataset papers find that primarily passive HFT strategies appear to have beneficial effects on market quality, such as by reducing spreads and reducing intraday volatility on average. Jovanovic and Menkveld (2012) find that the entry of a large, primarily passive HFT firm into the market for Dutch stocks was associated with a 15% decline in effective spreads. Malinova, Park and Riordan (2013) find that an increase in regulatory fees that primarily affected high frequency market making firms led to a significant increase in quoted and effective spreads in the Canadian equity markets. Hagstromer and Norden (2013) find that an increase in passive, market-making activity in the Swedish equity markets caused a decrease in short-term volatility.

In contrast, the HFT Dataset papers generally reveal that primarily aggressive HFT strategies raise more potential issues, with positive and negative aspects. On the positive side, aggressive HFT strategies can improve certain dimensions of price discovery, at least across very short time-frames. Brogaard, Hendershott and Riordan (2013) find that aggressive HFT activity improves price efficiency in the NASDAQ market by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors. Other papers reach more mixed findings on price discovery. Zhang and Riordan (2011) find that the information impact of HFT is significantly higher than non-HFT for large-cap stocks, but inconclusive for mid-cap stocks and significantly lower for small-cap stocks. Benos and Sagade (2013) find that aggressive HFT activity in the U.K. equity markets generates both significantly greater permanent price impact and significantly greater noise than non-HFTs. Zhang (2013) finds that aggressive HFT dominate price discovery in the short run (within a period of 10 seconds), but that passive non-HFTs consistently demonstrate a higher contribution to price discovery in the longer run (up to two minutes).

On the negative side, aggressive HFT activity also can impose costs on other market participants. Zhang and Riordan (2011) and Brogaard, Hendershott and Riordan (2013) find that aggressive HFT imposes increased adverse selection costs on non-HFT passive traders. Section III.B.3 below discusses two papers – Hirschey (2013) and Clark-Joseph (2013) – suggesting that some HFT firms may employ aggressive order anticipation and momentum ignition strategies that were highlighted in the Concept Release. Such strategies potentially can worsen the transaction costs of institutional investors and contribute to extreme volatility events.

Both Bershova and Rakhlin (2013) and Gao and Mizrach (2013) find that HFT is associated with increased intraday volatility, though they do not separately assess the effect of aggressive and passive HFT. Breckenfelder (2013) finds that competition among HFTs increases intraday volatility and reduces liquidity significantly, but has no effect on interday volatility. While he does not separately assess the effect of aggressive and passive HFT activity, he does measure the effect of competition on the extent to which HFTs engage in "liquidity consuming trades" (which can be aggressive or passive). He finds that the HFT ratio of liquidity consuming trades to total trades doubles from about 30% to 60% when HFTs compete for volume.

A few papers specifically examine the relation between HFT and the transaction costs of retail and institutional investors. Malinova, Park and Riordan (2012) find that an increased fee that primarily affected high frequency market making firms led to mixed results for retail traders – their aggregate transaction costs did not change, yet their intraday trading losses increased. In contrast, neither the aggregate transaction costs nor intraday returns of institutional traders were significantly affected by the fee change.

Tong (2013) finds that increased HFT activity on the NASDAQ market is associated with higher implementation shortfall costs incurred by institutional investors. In contrast, Bershova and Rakhlin (2013) conclude that HFT activity in the Tokyo and London equity markets is negatively correlated with the transaction costs of long-term investors. Neither Tong (2013) nor Bershova and Rakhlin (2013) separately examine the effects of aggressive and passive HFT activity.

Section III.B.4 below discusses papers that examine Flash Crash Datasets.[13] They provide insight into the conduct of HFTs during a severe trading disruption. Kirilenko et al. (2011) conclude that HFTs did not trigger the Flash Crash, but that their responses to the unusually large selling pressure that day exacerbated market volatility. As noted above, however, the analysis of HFT activity during the Flash Crash in Kirilenko et al. (2011) likely did not capture a large portion of HFT activity in the E-Mini that day.

With respect to equities, CFTC and SEC Staff (2010) find that HFTs continued to trade at high levels throughout the steep decline in broad index prices during the Flash Crash. Indeed, HFTs in the aggregate represented 50.3% of total volume at exchanges and other displayed venues during the rapid price decline in broad market prices. HFTs also were aggressive net sellers of $1.34 billion during that period. In contrast, HFTs in the aggregate represented only 36.6% of total displayed venue volume during the period when broad market prices recovered.

To conclude this summary, the staff believes it is important to highlight the limits on the scope of the economic literature to date in examining HFT. Due to the formidable data challenges facing researchers, the papers included in this literature review examine only a relatively small amount of HFT activity. The HFT Datasets generally have been limited to particular products or markets, and the data time periods now are relatively outdated, particularly given the pace of change in trading technology and practices. Accordingly, while the recent economic literature has made great progress in beginning to fill in the picture of HFT, much of the picture remains unfinished.

For example, the current literature does not reveal a great deal about the extent or effect of the HFT arbitrage strategies and structural strategies that were discussed in the Concept Release. Because the HFT Datasets generally have been limited to particular markets or products, they provide little opportunity to assess HFT strategies that simultaneously seek to capture price differentials across different products and markets. As further data on HFT activity becomes available, assessing these multi-product and multi-market strategies will be an important avenue of research.

III. Discussion of Papers

13 The events of the Flash Crash are described extensively in CFTC and SEC Staff (2010). These events will be discussed in a forthcoming part of this literature review that will address volatility more generally, including papers subsequent to CFTC and SEC Staff (2010) that shed further light on the dynamics of trading during the Flash Crash.

A. HFT Definitions and Factual Characteristics

As noted above, the HFT Dataset papers are significant because their authors have obtained access to non-public data that, to varying extents, can be used to identify HFT activity. An assessment of HFT Dataset papers must deal with the various metrics researchers used to define HFT and how their definitions may affect their conclusions about HFT activity. For example, the particular metrics used to classify HFT can greatly affect findings about key factual characteristics of HFT activity, including the extent of HFT volume as a percentage of total market volume and the division of HFT activity between aggressive trading and passive trading.

Given the importance of such threshold data and definitional issues, these issues are discussed first in this Section III.A, prior to discussing the relation between HFT and market quality in Section III.B below.

1. NASDAQ Datasets

Twelve papers examine NASDAQ Datasets, which are the only datasets available to academic researchers that directly classify HFT activity in U.S. equities.[14] The NASDAQ Datasets generally encompass the same broad ranges of data, but vary in certain particular respects.

The fullest description of a NASDAQ Dataset is provided in Brogaard, Hendershott and Riordan (2013). The NASDAQ Datasets include a stratified sample of 40 large-cap corporate stocks, 40 mid-cap corporate stocks, and 40 small-cap corporate stocks. Each market capitalization group is evenly split between NASDAQ listings and NYSE listings. The broadest trade information available in a NASDAQ Dataset includes 2008, 2009, and one week in 2010. The broadest quote information available in a NASDAQ Dataset includes NASDAQ Inside Quotes (the best displayed bids and offers on NASDAQ) available for the first week of each quarter during 2008 and 2009, the week of the Lehman bankruptcy in 2008, and one month in 2010.

NASDAQ used its access to trading and quoting activity on its market to identify the firms submitting orders. Based on its knowledge of the firms, NASDAQ manually classified 26 of the firms as HFT. The only stated factors for identifying the firms as HFT are how often their net trading in a day crosses zero (from long to short or short to long), the duration of their orders, and their order to trade ratio.

With the exception of Hirschey (2013), the papers that examine NASDAQ Datasets only have order and trade information that is aggregated across firms as "HFT" or "non-HFT." For each trade, the data identifies whether the aggressive (liquidity taking) side of the trade was HFT or non-HFT and whether the passive (liquidity providing) side of the trade was HFT or non-HFT. For NASDAQ Inside Quotes, the data indicates whether the order submitter was HFT or non-HFT.

[14] The NASDAQ Datasets papers are listed in Section V.A below.

Because HFT activity is aggregated across all 26 firms, the NASDAQ Datasets generally do not allow researchers to draw distinctions among different types of HFT firms. As discussed in Section III.B.3 below, Hirschey (2013) also has information tying activity to individual HFT firms, which he uses to measure persistence over time in the ability of particular HFT firms to profit from their aggressive orders.

Another limitation of the NASDAQ Datasets is that they do not cover all HFT activity. For example, they do not include HFT at firms that also act as brokers for customers because this activity cannot be clearly identified. The NASDAQ Datasets thereby exclude the proprietary trading desks of large integrated broker-dealer firms. The NASDAQ Datasets also do not include HFT at firms that route their orders through integrated firms because this activity cannot be clearly identified, which may exclude smaller HFT firms that rely on other firms for market access.

A third limitation of the NASDAQ Datasets is that they do not cover trading in corporate stock-related products, such as ETPs and equities-related futures. Accordingly, the NASDAQ Datasets are limited in the extent to which they can shed light on cross-product arbitrage strategies.

Despite the exclusions of certain types of HFT activity, the NASDAQ Datasets classify a large percentage of activity on NASDAQ as HFT. When measuring HFT activity, it is important to keep in mind the distinctions between different metrics of market activity. Every trade has two sides. Sides can be classified as buyer and seller or as aggressive (the side that traded immediately) and passive (the side that was resting on an order book when an aggressive order arrived). For purposes of this literature review, the term "HFT trade participation rate" means the extent to which HFTs participated on one or both sides of a trade. The term "HFT percentage of double-counted volume" means the percentage of total sides attributable to HFTs.

Carrion (2013) finds that, across his full sample, HFT firms had a trade participation rate of 68.3% of dollar volume, and that the HFT percentage of aggressive sides and passive sides were, respectively, 42.2% and 41.2% of dollar volume. These full sample measures of activity, however, mask very significant variations in activity across the three different size categories of stocks. For example, Brogaard, Hendershott and Riordan (2013) find that HFTs are much more active in large-cap stocks, and large-cap stocks represent the vast majority (96%) of total dollar volume in their sample. Accordingly, the full sample figures, which show a nearly even split between aggressive and passive trading, reflect almost entirely trading in large-cap stocks. For mid-cap stocks, in contrast, Brogaard, Hendershott and Riordan (2013) find that HFTs represented 36.3% of aggressive sides and 19.0% of passive sides. For small-cap stocks, HFTs represented 25.4% of aggressive sides and 11.2% of passive sides, resulting in a 69% aggressiveness ratio.[15] To summarize the paper's findings, HFTs are less active in mid-cap and small-cap stocks

[15] All aggressiveness ratios in this literature review will be calculated as the ratio of aggressive volume to total volume.

relative to large-cap stocks, and their activity in mid-cap and small-cap stocks is tilted much more toward aggressive trading rather than passive trading.

Gai, Yao and Ye (2013) examine a NASDAQ Dataset for October 2010 to assess the extent to which HFTs provide liquidity in stocks with low, medium, and high stock prices. Their focus is to determine whether competition on price among liquidity providers is constrained by the one-penny minimum tick increment, leading to a greater advantage for traders to compete on speed to achieve higher priority on exchanges with price-time priority rules.[16] For small-cap stocks, they find that HFTs provided approximately 19% of liquidity in all three price categories. For mid-cap stocks, they find that HFTs provided approximately 36%, 23%, and 22%, respectively, of liquidity in stocks with low, medium, and high prices. For large-cap stocks, they find that HFTs provided approximately 45%, 38%, and 31%, respectively, of liquidity in stocks with low, medium, and high prices. They interpret the increased percentage of HFT liquidity providers in stocks with lower prices as supporting a conclusion that HFTs focus their activity in stocks where competition on speed is more significant than competition on price.

Zhang and Riordan (2011) provide information on the percentage of quotes initiated by HFTs and their percentage of passive dollar volume (which generally occurs when a quote is matched with an aggressive order). They find that HFTs generate 73.7% of quotes, but execute only 43.7% of passive dollar volume.

O'Hara, Yao and Ye (2012) measure the influence of HFT on odd-lot trading. They find that 20-25% of trades initiated by HFTs are odd lots, and that trades initiated by HFTs are more likely to be odd lots than trades initiated by non-HFTs.

Several NASDAQ Dataset papers measure the trading profits of HFTs. Brogaard, Hendershott and Riordan (2013) measure the trading profits of aggressive trading (minus NASDAQ take fees) and passive trading (plus NASDAQ liquidity rebates).[17] They find that both aggressive trading and passive trading are profitable across all three market-cap categories, though passive trading would not be profitable across all three categories without liquidity rebates. They note that, if liquidity supply is competitive, then liquidity rebates should be incorporated into the endogenously determined spread (that is, liquidity suppliers consider the availability of rebates in determining the prices of their liquidity-supplying orders). Consistent with this competitive effect, they find that HFTs' liquidity supplying revenues are negative without fee rebates, which indicates that some of the rebates are being passed on to liquidity demanders in the form of tighter spreads.

Carrion (2013) also measures trading profits of HFTs. He finds that HFT passive trading is profitable even without liquidity rebates, while HFT aggressive trading is unprofitable

[16] The forthcoming Part III of this literature review will address aspects of Gai, Yao and Ye (2013) other than those that are related to their examination of a NASDAQ Dataset.

[17] NASDAQ charges take fees for aggressive orders that remove liquidity from its book and offers liquidity rebates for the passive orders against which aggressive orders are executed.

even before deducting NASDAQ taker fees. He also finds that both aggressive and passive HFTs engage in successful intraday market timing. The author explains that the difference in his results from Brogaard, Hendershott and Riordan (2013) is attributable to two factors – (1) he excludes trades where HFTs are both aggressive and passive; and (2) on days that end with an HFT inventory imbalance in a stock, he marks the imbalance with a volume-weighted average price of similar HFT trades in the stock, while Brogaard, Hendershott and Riordan (2013) mark the imbalance with the midpoint of the closing quotes in the stock.

2. E-Mini Datasets

Kirilenko, et al. (2011), Baron, Brogaard and Kirilenko (2012) and Clark-Joseph (2013) examine E-Mini Datasets. The E-Mini is the most actively traded instrument in the U.S. equities and equities-related futures markets and generally leads price discovery for the U.S. equity markets.[18] Unlike the U.S. equity markets, which are highly fragmented, the market for the E-Mini is fully centralized on a single exchange – the CME.

All three papers use transaction information in the E-Mini Datasets that identifies the individual trading accounts associated with the buyer and seller for the transaction, as well as whether the buyer or seller was aggressive or passive. A single firm may have multiple trading accounts.[19] In addition to transaction information, Clark-Joseph (2013) examines message data that includes the entry, cancellation, and modification of resting limit orders.

Unlike the NASDAQ Datasets, where NASDAQ manually classified firms as HFT, each of the E-Mini Dataset papers adopts a quantitative approach to classifying trading accounts as HFT. The approach to defining HFT varies for each paper, and the different results produced by the three approaches are instructive for understanding the nature of HFT, at least in a highly active instrument such as the E-Mini.

Kirilenko et al. (2011) was the first paper to examine an E-Mini Dataset and did so in the context of assessing E-Mini trading during the Flash Crash. The dataset includes transactions occurring between May 3, 2010 and May 6, 2010. To classify HFT trading accounts, the authors first define "Intermediaries" as accounts whose net holdings fluctuated within a quite narrow range (1.5% of end of day level), and whose end of day net position also was narrow (not more than 5% of its daily trading volume). They then classify as HFT those traders that ranked in the top 7% of Intermediaries in daily trading frequency. During the period from May 3-5, 2010, the defined HFT group encompassed 16 accounts that represented 34.2% of double-counted volume. On May 6, 2010, the

---

18 See CFTC and SEC Staff (2010) at 10.

19 Clark-Joseph (2013) notes that, throughout the E-Mini market, there exist assorted linkages between various trading accounts, such as where a single firm trades with multiple accounts. Accordingly, the trading account classifications in the E-Mini Dataset papers (such as "HFT," "opportunistic," and "fundamental") do not necessarily reflect all of the activity of a single proprietary trading firm. For example, a single proprietary trading firm may have multiple trading accounts, some of which are classified as HFT and some of which are not.

accounts classified as HFT represented 28.6% of double-counted volume. The majority of trading by the defined HFT group was passive, with an aggressiveness ratio of 46%.

Baron, Brogaard and Kirilenko (2012) adopt a more expansive approach to classifying accounts as HFT. They classify a trading account as HFT if it meets three criteria: (1) trading more than 5000 contracts per day (representing a notional value of more than $300 million in August 2010); (2) holding end of day inventory positions of no more than 5% of their total volume; and (3) maintaining intraday inventory positions of less than 10% of their total volume. During August 2010, the defined HFT group in Baron, Brogaard and Kirilenko (2012) encompassed 65 trading accounts that represented 54.4% of double-counted volume – as compared to the 34.2% estimated in Kirilenko et al. (2011).

Baron, Brogaard and Kirilenko (2012) further divided the 65 HFT trading accounts into "Aggressive" (at least 60% of their trades are liquidity taking), "Passive" (less than 20% of their trades are liquidity taking), and "Mixed." In August 2010, the 14 accounts classified as Aggressive HFT represented 15.2% of double-counted volume and executed 84.2% of their volume as aggressive. The 21 accounts classified as Passive HFT represented 8.9% of double-counted volume and executed only 12.4% of their volume as aggressive. The remaining 30 Mixed accounts represented 30.3% of double-counted volume and executed 37.1% of their volume as aggressive. The authors find that (1) the aggressiveness of a firm is highly persistent over the full two-year span, (2) the level of Aggressive HFT volume increases over the two-year period of the data sample from 15.2% to 22.6% of total volume, and (3) the profitability of Aggressive HFTs increased over that period, even as total E-Mini volume declined. Aggressive HFTs also were much more profitable than other HFTs. Aggressive HFTs earned a daily average of $45,267 in August 2010, while Mixed HFTs and Passive HFTs earned significantly less – respectively, $19,466 and $2,460.

Clark-Joseph (2013) adopts a third quantitative approach to classifying accounts as HFT. He first identifies each account whose end-of-day net position changed by less than 6% of its daily volume, and whose maximum intraday position changes are less than 20% of its daily volume. He then ranks the selected accounts by total trading volume and classifies the top 30 accounts as HFT. During the period between September 17, 2010 and November 1, 2010, the defined HFT group of 30 accounts represented 46.7% of double-counted volume and 48.5% of their volume resulted from aggressive orders. These 30 HFTs represented only 31.9% of all order entry, order modification, and order cancellation messages. The author further defined 8 HFT accounts as "A-HFTs" based on their profitability on aggressive orders, with the remaining 22 accounts classified as "B-HFTs." A-HFTs represented 24.7% of double-counted volume and executed 59.2% of their volume using aggressive orders. The B-HFTs represented 20.9% of double-counted volume and executed 35.9% of their volume using aggressive orders. The author also finds that the aggressive A-HFTs are more profitable than passive B-HFTs. Over the sample period, the A-HFTs earned a daily average of $99,168 per account, while the B-HFTs earn a daily average of $32,508 per account.

The varying approaches to classifying HFTs in the three E-Mini Dataset papers highlight several important points for evaluating HFT and assessing other papers. First, the precise metrics to define HFT can lead to widely differing results concerning the volume of trading that is classified as HFT. Second, there is great diversity among HFT trading accounts in terms of the aggressiveness of their trading. Some are highly aggressive and some are mostly passive. Moreover, the largest and most profitable HFT accounts are the ones that are most aggressive. Third, the first E-Mini paper – Kirilenko et al. (2011), which was in part the basis for examining E-Mini trading during the Flash Crash in CFTC and SEC Staff (2010) – adopted narrow parameters for classifying HFT that effectively excluded at least 20% of E-Mini volume from its classification of HFT. As a result, CFTC and SEC Staff (2010) accurately reflect the trading behavior of a subset of HFT activity in the E-Mini during the Flash Crash, but do not capture all of such HFT activity.

A limitation of the E-Mini Datasets is that they only encompass trading in one financial product. Consequently, as with the NASDAQ Dataset, the E-Mini Datasets are limited in the extent to which they can inform about HFT cross-market arbitrage and structural strategies. The absence of fragmentation in the E-Mini, however, allows for more exacting calculations of positions and profitability because traders cannot buy and sell the E-Mini in multiple trading venues.

3. Flash Crash Datasets

As discussed in Section III.A.2 above, Kirilenko et al. (2011) used an E-Mini Dataset to examine trading in the E-Mini during the Flash Crash, which was in part the basis for the analysis of E-Mini trading in CFTC and SEC Staff (2010). In addition, the discussion of equity market trading in CFTC and SEC Staff (2010) references two different datasets of HFT activity.

One dataset ("FINRA Dataset") consisted of the aggregate minute-by-minute dollar volume of trading on the day of the Flash Crash by the 12 largest HFTs, as reflected in audit trail data reported to FINRA. The FINRA Dataset included trades executed on NASDAQ and trades executed off-exchange and reported to NASDAQ's transaction reporting facility ("TRF"). FINRA manually classified firms as HFT. The 12 largest firms classified as HFTs had a trade participation rate of 46% on the day of the Flash Crash. The data did not allow HFT activity to be classified as aggressive or passive.

A second equities dataset ("Lit Venue Dataset") referenced in CFTC and SEC Staff (2010) consisted of all trades executed on the largest public quoting markets – each of the registered equities exchanges and the Direct Edge ECNs (which were not at that time registered as exchanges). In addition, specific participant data was obtained for the top 20 aggressive sellers on each venue during the Flash Crash. From this list, the staff manually identified 17 firms that appeared to be primarily associated with HFT. Similar to the NASDAQ Dataset, the Lit Venue Dataset did not include the proprietary trading desks of multi-service broker-dealers. The Lit Venue Dataset included trading volume in all securities across a 6-day period from May 3 through May 10, 2010. The data was aggregated by 15-minute increments, and HFT trade volume was classified as aggressive sells, aggressive buys, passive sells, and passive buys.

For the entire Lit Venue Dataset, the 17 HFT firms represented 43.8% of double-counted volume, with an aggressiveness ratio of 51.5%. In this respect, the Lit Venue Dataset, which encompassed trading in both corporate stocks and ETPs across all large lit venues, is generally consistent with the NASDAQ Datasets, which were limited to trading on NASDAQ in corporate stocks.

4. International Datasets

Fifteen papers examine HFT datasets relating to markets outside the U.S.[20] The International Dataset papers adopt varying approaches for classifying HFT activity that often result in important differences in the extent and type of activity that was examined.

a. Quantitative Classification of HFT

Four of the International Dataset papers use strictly quantitative metrics to classify trading activity as HFT or non-HFT (analogous to the E-Mini Dataset papers).

Malinova, Park and Riordan (2012) use non-public trader identifiers provided by the Toronto Stock Exchange to examine "intensive algorithmic trader" ("iAT) activity in the Canadian equity markets. The primary purpose of the paper was to assess the effect of an increased regulatory fee on messages, which includes orders, cancellations, and trades. To classify iATs, the authors required that a trader identifier be both in the top 5% of message-to-trade ratios and in the top 5% of the total number of messages submitted. As the authors noted, passive trading strategies generally require a much larger number of messages than aggressive strategies. Passive orders are left on a book awaiting contra-side aggressive orders and generally must be updated to reflect changing conditions. In contrast, aggressive traders choose their time of execution and price their orders to trade immediately. Using the two message-related metrics, the authors identify 94 trader identifiers as iAT. The classified iATs represented 53% of double-counted dollar volume, with only a 26% aggressiveness ratio. The authors note that their classification may capture high frequency market making by proprietary traders, as well as message-intensive agency algorithms that execute trading decisions on behalf of an institutional client.

Similarly, IIROC (2012) adopts a quantitative approach for classifying HFT in the Canadian equity markets that focuses on traders with the highest order-to-trade ratios. The approach classifies 11% of traders on the Canadian equity markets as HFTs. These traders had an average order-to-trade ratio of 11.2 to 1 and higher. They represented 32% of double-counted dollar volume, yet also accounted for 94% of order messages. The classified HFTs executed 66% of their volume passively during the sample period of August through October of 2011.

[20] The fifteen International Dataset papers are listed in Section V.C below.

ASIC (2013), in turn, adopted a quantitative approach for classifying HFT activity that relies on six metrics – (1) order-to-trade ratios, (2) the number of "fast messages" (number of posted messages within 40 milliseconds of a defined event), (3) holding time (volume-weighted time that a position is held), (4) end-of-day positions as a percentage of a firm's total volume, (5) total volume per day, and (6) at-best ratio (number of orders placed at the best posted price plus the number of orders priced at market, divided by the total number of orders). Traders are scored cumulatively across all six metrics, with the highest 15% on any given day classified as HFT. The approach yields a group of HFTs that represented 27% of double-counted volume during the sample period of May through July of 2012, with an aggressiveness ratio of 46.5% for their continuous market volume.[21]

In contrast, Kang and Shin (2012) adopt a quantitative approach to HFT classification that captured primarily aggressive traders as HFT. The authors examine trading in KOSPI 200 futures on the Korea Stock Exchange. With trader identifiers provided by the exchange, they use three metrics to classify traders as HFT – (1) number of daily messages of more than 1000 (approximately one message every twenty seconds during a normal trading day); (2) an end-of-day inventory position of less than 3% of its trading volume; and (3) a median cancellation time for cancelled limit orders of less than two seconds. Importantly, this third factor was not determined by number of cancelled limit orders, but by how quickly limit orders were cancelled. This approach led to classifying an average of only four trading accounts per day as HFT, but they represented 15% of double-counted volume, with an aggressiveness ratio of 74%.

In sum, the International Dataset papers that employ quantitative approaches to classifying HFT reached widely varying findings regarding the factual characteristics of HFT. Malinova, Park and Riordan (2012) and IIROC (2012) focused exclusively on message rates, which resulted in primarily passive trading being classified as iAT. Kang and Shin (2012), in contrast, focused on how quickly limited orders are cancelled, which led to classifying primarily aggressive strategies as HFT. ASIC (2013) used a six-metric approach that yielded a more even split between aggressive and passive trading.

b. Manual Classification of HFT

Other International Dataset papers use manual approaches to classify trading activity as HFT or non-HFT (analogous to NASDAQ's approach to classifying HFT firms for the NASDAQ Datasets).

Hagstromer and Norden (2013) use trader identifiers and knowledge of member activities provided by NASDAQ OMS Stockholm ("OMXS") to manually classify firms as HFT. Their sample was limited to the actively traded stocks in the OMXS 30 and included two

[21] In total, the HFTs in ASIC (2013) executed 46% passively in continuous trading, 40% aggressively in continuous trading, and 14% in auctions or crossings. Other HFT Dataset papers generally focus their analysis of HFT on continuous trading and exclude auctions and crossings.

months of trading – August 2011 and February 2012. The authors classify traders as HFT if they engaged in proprietary trading only and use algorithms in their trading strategies. This approach resulted in 29 traders being classified as HFT. Overall, HFTs represented 25-29% of double-counted volume, of which 57-59% was executed passively.

Hagstromer and Norden (2013) further divide classified HFTs into "market-making" and "opportunistic." The authors classify HFTs as market-making if their average frequency of posting limit orders at either side of the inside quotes was more than 20%. Market-making HFTs represented 65-71% of HFT volume, with an aggressiveness ratio of 29-32%. The remaining HFTs were classified as opportunistic. Opportunistic HFTs represented 29-35% of HFT volume, with an aggressiveness ratio of 64-68%. Notably, both the percentage of HFT volume and the aggressiveness ratios of opportunistic HFTs increased substantially from August 2011 to February 2012 – respectively, from 29% to 35% of HFT volume and from 64% to 68% aggressiveness ratios. The level of order submissions also differed significantly between market-making HFTs and opportunistic HFTs. For example, market making HFTs represented 65% of HFT volume and 86% of HFT orders submissions in February 2012, while opportunistic HFTs represented 35% of HFT volume and 14% of HFT order submissions. These results highlight the tilt in HFT classification toward passive trading that can result from focusing on the level of order submissions as a metric of HFT.

In addition, Hagstromer and Norden (2013) examine changes in HFT that are associated with tick size changes. On the OMXS, tick sizes are determined by stock price level, with wider tick sizes kicking in at specified higher price levels. They find that tick size increases have a strongly negative effect on the market share of aggressive HFTs, while tick size increases have a positive effect on the market share of passive HFTs.

Hagstromer, Norden and Zhang (2013) adopt a similar methodology as Hagstromer and Norden (2013) to assess the use of order types by different groups of traders. With respect to order cancellation rates, they find very small economic differences between HFTs and non-HFTs.

Brogaard, Hagstromer et al. (2013) examine trading on the OMXS with a dataset that includes member identifiers, as well as information on whether members purchased varying levels of co-location services offering faster trading speeds. Among other things, they apply the definitions of HFT used in Hagstromer and Norden (2013) and Kirilenko et al. (2011) to their dataset. They find that use of co-location status as an indicator of trading speed differs substantially from the two HFT definitions. For example, they find that numerous non-HFTs (according to both definitions) use co-location accounts. In addition, the Kirilenko et al. (2011) definition of HFT identified 9 accounts with 37.7% of double-counted trading volume as HFT, while the Hagstromer and Norden (2013) definition identified 64 accounts with 33.1% of double-counted trading volume as HFT. For the Hagstromer and Norden (2013) definition of HFT, 24 HFT accounts with 25.8%

of double-counted trading volume were co-located, while 40 HFT accounts with 7.3% of double-counted trading volume were not co-located.[22]

Breckenfelder (2013) uses access to trader identifiers to examine trading on the OMSX in OMXS 30 stocks. He classified fewer than ten firms (the exact number was not specified for confidentiality reasons) as large international HFTs. He does not provide statistics for the group as a whole, but does provide statistics for the two HFTs he determines to be most different from each other. He finds that a key characteristic on which the two HFTs differ is aggressiveness – one HFT ("HFT A") executed 91% of its trades aggressively, while the other ("HFT B") executed only 35% of its trades aggressively. The author also develops a metric of "liquidity consuming trades" – defined as a stock buy (whether aggressive or passive) after a price increase and a stock sell (whether aggressive or passive) after a price decline. He characterizes this metric as capturing the extent to which HFTs trade with the price trend, whether aggressively or passively. With this metric, he finds much less difference between the two firms – HFT A has a ratio of liquidity consuming trades to total trades of 60%, while the ratio for HFT B is 54%.

Benos and Sagade (2013) use U.K. regulatory transaction data to examine HFT activity in four randomly selected FTSE 100 stocks over a randomly selected one-week period in 2010-2012 (the stocks and dates were not revealed for confidentiality reasons). The data cover trades executed on four large lit venues – the London Stock Exchange ("LSE"), Chi-X, BATS, and Turquoise. The authors manually classify firms as HFTs using press reports and company websites to see whether they described themselves by such terms as "HFT," "low-latency trader," or "electronic market maker." The firms classified as HFT represented 27.1% of double-counted volume, with an aggressiveness ratio of 46.9%. The authors then divide the HFTs into an aggressive group with above-median aggressiveness and a passive group with below median aggressiveness. The Aggressive HFTs represented 16.3% of double-counted share volume, with a 56.3% aggressiveness ratio, while the passive HFTs represented 10.8% of double-counted share volume, with a 32.7% aggressiveness ratio.

Benos and Sagade (2013) highlight that the granular nature of their dataset enables them to see a wide diversity of trading strategies among HFT firms, with aggressiveness ratios ranging from 30% to 75%. They emphasize that, when examining the effect of HFT activity on market quality, different types of strategies need to be examined separately. They found, for example, that Passive HFTs follow a trading strategy consistent with market making and, as such, their trades have alternating signs (buy and sell) and are independent of recent (ten-second) price changes. By contrast, they find that aggressive HFTs exhibit persistence in the direction of their order flow and trade in line with the recent (ten-second) price trend. Finally, they find that both groups of HFTs are sensitive to inventory imbalances over a longer time period, which they note is consistent with the idea that HFTs try to end the day with a relatively flat position.

---

[22] Due to the small number of firms classified as HFT using the Kirilenko et al. (2011) definition, confidentiality restrictions prevented Brogaard, Hagstromer et al. (2013) from disclosing the distribution of HFTs across co-location and non-co-location accounts.

Alampieski and Lepone (2011) and (2012) use U.K. regulatory data to examine HFT activity in FTSE 100 stocks for 30 trading days in 2010. The data covered trading on three lit venues – LSE, Chi-X, and BATS. The U.K. regulator and the trading venues manually identified traders as HFT based on their use of "computer algorithms and low latency infrastructure to generate and execute trading decisions for the purpose of generating returns on proprietary capital." They identified 363 traders associated with 52 unique firms as HFT. Similar to the NASDAQ Dataset, trade and order information is aggregated as "HFT" or "non-HFT" so that there was no visibility into individual trader strategies. The primary purpose of Alampieski and Lepone (2011) is to compare HFT activity on the days on which U.S. macroeconomic data was announced to non-announcement days. The authors find that the HFT share of aggressive trading volume ranged from approximately 20% to 27.1% on non-announcement days and from 20% to 34.8% on announcement days. The HFT share of passive trading volume ranged from approximately 25% to 27.4% on non-announcement days and from approximately 25% to 37.0% on announcement days.

Brogaard, Hendershott et al. (2012) use U.K. regulatory data to examine trading in FTSE 200 stocks. Employing a definition of HFT similar to the one in Alampieski and Lepone (2011), the three major trading venues for U.K. stocks – the LSE, BATS, and Chi-X – used their understanding of the business of their participants to identify 52 as HFTs. The HFT percentage of double-counted volume reached a high of approximately 23% in 2010. The authors did not provide statistics on the HFT share of aggressive and passive trading.

Two other papers, Jovanovic and Menkveld (2012) and Menkveld (2013), examine the trading of a single HFT firm. The authors used anonymized broker identifications to manually classify one large firm as HFT, primarily because of its large percentage of trading in Dutch stocks on Chi-X, a new venue for European securities. The firm also executed a large volume (though smaller percentage) of trading on the Euronext exchanges. The HFT firm participated in 14.4% of all trades in the dataset (64.4% of Chi-X trades and 8.1% of Euronext trades). In both markets, the single HFT firm had an aggressiveness ratio of only 22%.

Finally, Bershova and Rakhlin (2013) use a dataset on trading in the Tokyo and London equity markets that was obtained from a large, multi-service broker with a significant presence in both markets during the first two quarters of 2010. The authors classified HFT clients as those that use the broker's ultra-low latency infrastructure. The HFT clients of this single broker represented approximately 11% of volume in TOPIX stocks and 5% of volume in FTSE 350 stocks.[23] The dataset did not permit volume to be classified as aggressive or passive.

B. Relation Between HFT and Market Quality

[23] It is not clear whether the HFT percentages in Bershova and Rakhlin (2013) are of single-counted or double-counted volume.

The HFT Dataset papers examine the relation between HFT and a variety of aspects of market quality. They can be divided into four categories: (1) papers that examine more general aspects of market quality, particularly spreads, price discovery, volatility, and liquidity; (2) papers that focus on the transaction costs of retail and institutional investors; (3) papers that address the two potentially problematic strategies highlighted in the Concept Release – order anticipation and momentum ignition, and (4) papers that focus on HFT during a severe market disruption – the Flash Crash.

1. Spreads, Price Discovery, Volatility, and Liquidity

HFT Dataset papers that focus on the relation between HFT and different types of spreads, [24] various dimensions of price discovery, short-term volatility, and liquidity generally indicate that these effects vary significantly depending on whether HFTs engage in aggressive activity or passive activity.

Jovanovic and Menkveld (2012), for example, find that the entry of a large, primarily passive HFT firm into the market for Dutch stocks was associated with a 15% decline in effective spreads and a 23% reduction in the adverse selection costs of posted prices.

Bershova and Rakhlin (2013) examine data on aggregate HFT and long-term institutional investor trading volume that was routed through a single, large multi-service broker in the Tokyo and London equity markets.[25] They find that HFT is associated with a significant compression of bid-ask spreads, but also with an increase in short-term volatility.

Malinova, Riordan and Park (2013) find that quoted and effective spreads increased significantly with the decline in primarily passive iAT message rates following a

---

[24] Zhang and Riordan (2011) measure quoted spreads, effective spreads, realized spreads, and price impact. Quoted spread is an ex-ante measure of liquidity and is calculated as the difference between the best (highest) bid and best (lowest) offer at trade time. Effective spread is an ex-post measure of liquidity and is calculated as the difference between the trade execution price and the midpoint of the best bid and offer at trade time. An effective spread can be decomposed into realized spread and price impact. Realized spread is calculated as the difference between the trade execution price and the midpoint of the best bid and offer at a specified time interval after trade time (the authors use 5 and 15 minutes). Price impact is calculated as the difference between the midpoint of the best bid and offer at trade time and the midpoint at a specified time interval after trade time (the authors use 5 minutes and 15 minutes). Stated another way, realized spread measures the potential for a liquidity supplier to profit from a trade by liquidating the position at some specified point in the future. If price impact (an adverse price move from the standpoint of the liquidity supplier) exceeds the effective spread, the realized spread will be negative. A negative realized spread means the liquidity supplier potentially incurred a short-term trading loss due to an adverse price move after the trade. When realized spreads are categorized by the type of counterparty that is taking liquidity, this measure can capture whether the type of counterparty is informed about short-term price changes.

[25] Bershova and Rakhlin (2013) define HFT firms as users of the broker's ultra-low latency infrastructure, while long-term investors are defined as users of the broker's traditional direct market access and strategy access, smart order routers, and dark pools.

regulatory fee increase, and that realized spreads (which measure the potential revenue or losses of liquidity suppliers) decreased significantly.

Zhang and Riordan (2011) use a NASDAQ Dataset to measure differences in spreads and information impact when HFTs and non-HFTs are trading passively and aggressively. In general, they find that HFTs tend to trade passively when spreads are wide and trade aggressively when spreads are narrow. They also find that realized spreads are negative when non-HFTs are liquidity suppliers, and that negative realized spreads are significantly greater for trades when HFTs are the aggressive trader than when non-HFTs are the aggressive trader. The authors further find that the information impact (as measured by Hasbrouck impulse response functions) of HFT is significantly higher than non-HFT for large-cap stocks, but inconclusive for mid-cap stocks and significantly lower for small-cap stocks. They state that "we might infer that HFT bring information into the market only for high market capitalization stocks, although we cannot distinguish whether the information impact results from the speed advantage due to information processing algorithms, connections speed, or whether HFT use order anticipation methods."

Carrion (2013) also finds that HFTs tend to be more aggressive when spreads are narrow and more passive when spreads are wide, and that realized spreads are significantly more negative when HFTs trade aggressively. He further examines the effect of HFT activity on price efficiency. He measures price efficiency in two ways: (1) order imbalance – the incorporation of information from lagged order flows within 1 minute and 5 minute periods; and (2) price delay – the incorporation of information from market index returns. He finds that price efficiency as measured by both metrics is significantly higher on days of high HFT aggressiveness.

Brogaard, Hendershott and Riordan (2013) use a NASDAQ Dataset to examine the effect of HFT activity on short-term price efficiency. They assess price efficiency with a state space model that measures the effect of HFT trading on the permanent and transitory components of price impacts.[26] They find that HFT aggressive trades facilitate price efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors, both on average and on the highest volatility days.

Brogaard, Hendershott and Riordan (2013) further find that HFT aggressive trading imposes adverse selection costs on non-HFT passive traders. Given that HFTs predict price changes occurring only a few seconds in the future, they conclude that the short-lived nature of HFTs' information raises questions about whether the price efficiency gains outweigh the direct and indirect adverse selection costs imposed on non-HFTs. They note, for example, that indirect adverse selection costs could include increases in

[26] Brogaard, Hendershott and Riordan (2013) describe their state space model as assuming that a stock's price impact can be decomposed into a permanent component and a transitory component. The permanent (efficient) component is the part of price impact due to information about a stock's fundamentals, while the transitory component of price impact results from a temporary liquidity imbalance. Papers may refer to transitory price impact as pricing error or noise.

technology costs by non-HFTs to increase their speed of trading to avoid adverse selection.

Finally, Brogaard, Hendershott and Riordan (2013) find that passive HFT trades are adversely selected (subsequent prices move against the position – down for buys and up for sells), as well as negatively associated with permanent price impact and positively associated with transitory price impact. Similarly, Menkveld (2013) finds that the primarily passive HFT firm examined in his paper incurs adverse selection costs and that its trading is positively associated with pricing errors (transitory price impact). He notes that this effect of the HFT's passive trading on transitory pricing errors largely supports the market-making character of the HFT's activity and is consistent with an interpretation that the primarily passive HFT acts as a modern market maker.

Benos and Sagade (2013) assess the impact on market quality of primarily aggressive HFTs and primarily passive HFTs. In particular, they use the Hasbrouck vector autoregressive ("VAR") framework coupled with a variance decomposition technique to measure the effect of each group's aggressive trades on permanent price impact and noise.[27] They find that the aggressive trades of each group generate both greater permanent price impact and greater noise than their volume participation, though the primarily aggressive HFT group generates more permanent price impact and noise than the primarily passive HFT group. The authors note that it appears that the more HFTs trade aggressively, the more they contribute to both price discovery (as measured by permanent price impact) and excess volatility (as measured by noise).

Benos and Sagade (2013) distinguish their findings regarding price efficiency from those of Brogaard, Hendershott and Riordan (2013) by emphasizing the granular nature of their dataset (which allowed them to tie trading to individual HFT accounts) compared to the aggregated nature of the NASDAQ Dataset (in which HFT activity was aggregated across 26 firms). Their dataset allowed them to examine separately the impact of aggressive trades by primarily aggressive HFTs and the impact of aggressive trades by primarily passive HFTs. In contrast, the aggregated nature of the NASDAQ Dataset means that aggressive trades of both primarily aggressive HFTs and primarily passive HFTs were lumped together in Brogaard, Hendershott and Riordan (2013).

Benos and Sagade (2013) conclude that the overall welfare implications of HFT activity are unclear. They state that such implications will depend on how the marginal benefit of

[27] The Hasbrouck VAR framework is a method researchers use to separate the permanent and transitory components of price impact. Benos and Sagade (2013) state that this method clarifies the relationship between price efficiency and volatility. They note, for example, that if prices are more volatile in the presence of HFT because they incorporate new information about fundamentals, then HFTs make a positive contribution to price efficiency. If, on the other hand, HFTs cause prices to move away from fundamentals, then the resulting volatility is noise and HFT is detrimental to market quality. The authors state that the Hasbrouck VAR framework empirically distinguishes between these two effects by assuming that prices have efficient (non-stationary) and noisy (stationary) price components, and estimates the contribution of HFTs to the variance of each component.

information at some times compares with the marginal cost of excess volatility at other times, including in periods of market stress.

Zhang (2013) uses a NASDAQ Dataset to examine the role of HFTs in reacting to two kinds of information shocks: (1) "hard" information shocks as measured by the top 1% of ten-second price changes in the E-Mini and the CBOE Volatility Index ("VIX"); and (2) "soft" information shocks as measured by positive and negative firm-specific newswire items.[28] She finds that aggressive HFTs strongly react to hard information shocks, particularly changes in E-Mini prices (aggressive HFT responses to E-Mini price shocks were nearly 10 times greater than such responses to VIX price shocks). Aggressive HFTs traded with the direction of the E-Mini shock in the first 10 seconds, but completely reversed their position within two minutes after the shock. She interprets these findings as a realization of HFT short-term profits. While aggressive non-HFTs also increase their trading in the direction of the E-Mini shock in the first 10 seconds, they do not reverse their positions within 2 minutes. She finds opposite results when considering passive activity. Both HFTs and non-HFTs suffer adverse selection in the first ten seconds by net trading against the price move, but passive HFTs are much less selected against than passive non-HFTs.

Zhang (2013) further finds that HFTs respond quite differently to soft information shocks. Although both aggressive HFTs and non-HFTs trade in the direction of soft information shocks, non-HFTs demonstrated a significantly stronger reaction (approximately 3 times stronger net trading) than the HFT reaction. Also, in contrast to the results for E-Mini shocks, both passive HFTs and passive non-HFTs were adversely selected to a similar extent following soft information shocks.

To assess the robustness of her results, Zhang (2013) separately measures the reactions to price shocks before, during, and after the financial crisis of 2008-2009 (measured as periods when the VIX was under 30, above 30, and then back under 30). She finds that aggressive HFT responses to E-Mini price shocks during the crisis were nearly four times higher than before the crisis and that post-crisis responses remained more than three times higher than pre-crisis responses.

In addition, Zhang (2013) assesses the effect of HFT and non-HFTs on price discovery by testing whether HFT order flow has a significant influence on stock returns relative to non-HFT order flow. She finds that aggressive HFT orders dominate price discovery in the short run (within a period of 10 seconds). In the longer run (up to two minutes), however, she finds that passive non-HFTs consistently demonstrate a higher contribution to price discovery than HFTs.

O'Hara, Yao and Ye (2012) assess the effect of HFT and non-HFT odd-lot trades in a NASDAQ Dataset on daily price discovery. They find that odd-lot trades between two non-HFT participants play a much larger role in daily price discovery than any category

[28] Zhang (2013) characterizes hard information as quantitative and easily processable and storable by computers, while soft information is qualitative and hard to interpret by computers.

of trade involving an HFT. Johnson, Van Ness and Van Ness (2013) also assess the effect of HFT and non-HFT odd-lot trades in a NASDAQ Dataset by using weighted price contribution as a measure of price discovery. They find that weighted price contribution is positive for odd-lot trades when a non-HFT firm provides liquidity, while weighted price contribution is negative for odd-lot trades when an HFT firm provides liquidity.

Hagstromer and Norden (2013) use a natural experiment based on tick size changes to assess the effect of passive HFT activity on short-term volatility. On the OMXS, tick sizes are determined by stock price level, with wider tick sizes kicking in at specified higher price levels. The authors assess the relationship between tick size increases and volatility with an event study that divides trading into two time periods – (1) a period before passive HFTs began trading on OMXS and when HFTs were primarily aggressive, and (2) a period after passive HFTs began trading on OMXS in addition to aggressive HFTs. While tick size increases during the first period were associated with increases in volatility, tick size increases during the second period resulted in no significant volatility changes. The authors conclude that an increase in passive market-making HFT activity causes a decrease in short-term volatility.

Breckenfelder (2013) examines OMXS trading to assess the effect of HFT on short-term volatility and liquidity. His focus is on competition between HFT firms. Accordingly, he examines changes in volatility and liquidity when one HFT firm begins or ceases competing with another HFT firm. He finds that competition among HFTs increases intraday five-minute volatility by 9%, hourly volatility by 20%, and maximum intraday volatility by 14%, but has no effect on interday volatility. He also finds that liquidity, as measured by the Amihud measure of illiquidity,[29] decreases significantly in the presence of HFT competition. Five-minute liquidity decreases by 9%, and 60-minute liquidity decreases by 15%.

While Breckenfelder (2013) does not separately measure the effects of aggressive HFT and passive HFT on volatility, he does measure the effect of competition on the extent to which HFTs engage in liquidity consuming trades (which, as defined in Section III.A.4.b above, can be both aggressive and passive). He finds that the HFT ratio of liquidity consuming trades to total trades double from about 30% to 60% when HFTs compete for volume.

Breckenfelder (2013) also checks the robustness of his main analysis by examining the effect of a change in tick size regime on HFT competition and market quality. In 2009, a tick size harmonization by the Federation of European Securities Exchanges ("FESE") led to a general reduction in tick sizes for most OMXS stocks. He divides his sample stocks into one group for which the tick size did not change, another group for which tick size decreased without any change in HFT competition, and another group for which tick

[29] The Amihud measure of illiquidity essentially captures the extent to which price changes are associated with trading volume – the lower the volume relative to a price change, or the larger the price change relative to volume, the higher is the illiquidity.

size decreased and HFT competition increased. For both groups in which tick sizes decreased, he finds that HFT activity increased – from 8% to 16% for the second group and from 8% to 24% in the third group. He finds that the decrease in tick size without any change in HFT competition led to an increase in market quality, although its statistical significance was low or absent. In contrast, a decrease in tick size with an increase in HFT competition led to a statistically significant decrease in market quality (as measured by intraday volatility and liquidity) consistent with his main analysis.

Gao and Mizrach (2013) use a NASDAQ Dataset to examine the effect of HFT on "market quality breakdowns," which they define as an intraday decline in a symbol of at least 10% below its 9:35 a.m. price, followed by an intraday recovery to at least 2.5% of 9:35 a.m. price. The authors generally find that the frequency of these periods of extreme intraday volatility has fallen in recent years.[30] They also examine whether HFT, as identified in the NASDAQ Dataset, is associated with market quality breakdowns. They find that HFT causes correlation among market orders in different stocks, and that this correlation is associated with an increase in breakdown frequency of 31%. They also find that spikes in HFT activity raise the breakdown frequency by an additional 18.3%.

Gerig (2012) uses a NASDAQ Dataset to examine the role of HFT in synchronizing prices among stocks. For 35 of the large-cap stocks in the NASDAQ Dataset, he first uses historical transaction data to show that the time period for synchronization (the average normalized price response of a security to a price movement in another security) has fallen sharply. This synchronization took several minutes in 2000, approximately 1 minute in 2005, and less than 10 seconds in 2010. He then uses the NASDAQ Dataset to separate the synchronization response in 2010 between HFT activity, non-HFT activity, and activity that could not be categorized. He finds that the overwhelming majority of the initial price response is due to HFT activity. He also examines the extent to which synchronization is associated with HFT activity in each of the 120 stocks in the NASDAQ Dataset. He finds that HFT activity varies significantly from security to security, and that synchronization explains the majority of the variance. He concludes that synchronization can benefit average investors by diffusing information rapidly from security to security, but pricing errors potentially may propagate quickly throughout the financial system and lead to market instability during times of stress if safeguards are not in place.

Finally, ASIC (2013) conclude that HFT in the Australian equity markets does not appear to be a key driver for changes seen in price formation, liquidity, and execution costs, and that HFT does not exacerbate market instability. They further find that the contribution of HFT passive orders to depth around the best prices remained relatively constant in periods of high volatility, but also found that HFTs increased their aggressive trading during these periods.

[30] This general finding of Gao and Mizrach (2013) on market quality breakdowns will be discussed further in a forthcoming part of this literature review that will cover market volatility, among other things.

2. Transaction Costs of Retail and Institutional Investors

The authors of four papers obtained trading data with identifiers for HFT activity and for the activity of retail or institutional investors. The additional datasets enable a more direct assessment of the effects of HFT on the transaction costs of these investor groups than is possible with the more general market quality metrics discussed in Section III.B.1 above.

Malinova, Park and Riordan (2013) use a regulatory dataset with trader identifiers and a proprietary dataset of retail trader transactions to examine the effect of iATs (intensive algorithmic traders – which they note include high frequency market makers) on retail traders and institutional traders in the Canadian equity markets. They focus on market activity after the introduction of a regulatory fee per exchange message. The application of the fee was initially unknown and had the potential to be very costly for high message traffic participants. As noted in Section III.A.4 above, the authors defined iATs in terms of message rates. This approach produced a sample of iAT activity with only a 26% aggressiveness ratio.

The authors use the retail trader dataset to directly measure the effect of the fee change on the trading costs and intraday trading profitability of retail traders. They also use their access to trader identifiers to measure the effect of the fee change on institutional traders. Institutional traders are defined as those that are neither iATs nor retail traders and that had an absolute cumulative net position that exceeded $25 million in at least one stock. The $25 million bound corresponded to selecting approximately the top 5% of identifiers with regards to maximum net position.

After the fee change, iATs reduced their total messages by over 30%. The authors find that quoted and effective spreads rose significantly with the decline in iAT message rates and that realized spreads (a measure of profitability for passive traders) decreased significantly. The authors suggest that, after the fee change, passive traders were hampered by their inability to manage their market exposure through limit order cancellations.

The effect of the fee change on retail traders and institutional traders was mixed. First, despite the increase in quoted and effective spreads, the aggregate trading costs of both groups did not change significantly. The authors note that their measure of aggregate trading costs included the ability to earn spreads on passive orders, and that the increase in spread costs for aggressive orders after the fee change was offset by a reduction in costs attributable to passive orders. As a result, retail and institutional traders incurred increased effective spread costs for their aggressive orders, but earned increased effective spreads for their passive orders. The net result was that their aggregate trading costs did not change significantly, despite the increase in quoted and effective spreads that followed the fee change.

Turning to intraday returns of traders, the returns of institutional traders did not change significantly, while the intraday trading losses of retail investors increased significantly after the fee change. Intraday trading losses were measured either as actual intraday losses or, if there was an inventory position at end of day, using closing prices. Retail traders incurred intraday losses both before and after the fee change, but the intraday losses increased significantly after the fee change.

Three papers focus on the effects of HFT activity on the trading costs of institutional investors that, unlike retail traders, typically need to trade in large size.

Tong (2013) examines a NASDAQ Dataset and an Ancerno Ltd. proprietary database of institutional investors' equity transactions ("Ancerno Dataset"). She finds that increased HFT activity across all three capitalization groups of stocks – small, medium, and large – is associated with higher implementation shortfall costs incurred by institutional investors.[31] In particular, she finds that, after controlling for other economic determinants of execution shortfall, a one standard deviation increase in HFT activity leads to a 5 basis point increase in execution shortfall. Moreover, she finds that the effect is most pronounced in smaller stocks. She does not separately test the effect of aggressive HFT and passive HFT on institutional trading costs.

Brogaard, Hendershott et al. (2012) assess the effect of HFT on institutional trading costs in U.K. equity markets. They use a regulatory dataset of HFT activity and Ancerno data containing institutional investor orders and transactions, but employ a methodology that yields different results than Tong (2013). They focus on institutional investor transaction costs following two LSE technology upgrades that reduced latency by 1.0 millisecond and 0.7 milliseconds. The two upgrades were followed by increases in HFT activity that ranged from 2 to 7 percentage points. Interestingly, two other upgrades that reduced latency by 5.0 milliseconds and 1.0 millisecond were not followed by an increase in HFT activity. For the two upgrades that did increase HFT activity, the authors focus on 10-day periods before and after the upgrades. The authors do not find a statistically significant relationship between the technology upgrades and institutional investor transaction costs following the two upgrades.

Finally, Bershova and Rakhlin (2013) examine data on aggregate HFT and long-term institutional investor trading volume in the Tokyo and London equity markets that was routed through a single, large multi-service broker. Long-term investors were defined as users of traditional direct market/strategy access services, smart order routers, and dark pools. As noted in Section III.A.4.b above, the authors classified HFT clients as users of the broker's ultra-low latency infrastructure. They find that HFT is negatively correlated with the transaction costs of long-term investors. As noted above in Section III.B.1, they further find that HFT is associated with a significant compression of bid-ask spreads, but also with an increase in short-term volatility. They conclude that the transaction costs for long-term investors associated with the increase in volatility were more than offset by the reduction in bid-ask spreads. The authors emphasize, however, that the long-term

[31] Tong (2013) defines "execution shortfall" (also known as implementation shortfall) as the percentage difference between the average execution price of an order and a benchmark price that is prevailing in the market when the order ticket is placed with a broker.

investor orders in their dataset reflected primarily liquid stocks with a flat daily close-to-open profile, where passive liquidity provisioning is a dominant strategy for electronic market makers. They therefore caution that their findings should not be generalized to all long-term investor orders. For example, they expect that the reduction in spreads would not offset the effect of higher short-term volatility for large orders.

3. Order Anticipation and Momentum Ignition Strategies

As noted in Section I.B above, the Concept Release requested comment on two types of strategies that may pose particular problems for long-term investors – order anticipation and momentum ignition. Two papers find evidence suggesting that HFTs employ these strategies when trading aggressively.

Hirschey (2013) uses a NASDAQ Dataset to examine whether HFT aggressive trades could increase non-HFT's trading costs by anticipating and trading ahead of non-HFT order flow. For example, he notes that HFTs may anticipate the trades of a mutual fund if the mutual fund splits large orders into a series of smaller ones and the initial trades reveal information about the mutual fund's future trading intentions. He finds that HFT aggressive trades lead those of non-HFTs. When HFTs sell stock aggressively during a particular second, this aggressive selling forecasts future aggressive selling by non-HFTs that continues through five minutes into the future. He finds this effect in all three capitalization groups in the NASDAQ Dataset, and it is strongest in small-cap stocks. The author also finds that the anticipatory trading effect is stronger at times when non-HFTs may be more impatient, such as on high volume and high buy-sell order imbalance days. He explores whether these results are explained by HFT's reacting faster to news, by positive-feedback trading by non-HFTs, or by HFTs and non-HFTs trading on the same signals. He concludes, however, that the results are best explained by HFT anticipatory strategies.

In addition, Hirschey (2013) has access to information on individual HFT firm identifiers in the NASDAQ Dataset (in addition to the aggregated HFT information used in other papers). He finds that HFTs vary in their skill at predicting non-HFT order flow, and that aggressive trades by the HFTs that are most skilled at predicting non-HFT order flow also predict larger price changes than do the aggressive trades of other HFTs. He concludes that these findings are consistent with the significant use of order anticipation strategies by some HFTs. He notes that, if a non-HFT is informed about fundamental values, then the anticipatory strategies of HFTs would capture some of the non-HFT's information rent and, as a result, undermine their incentives to do fundamental research. He notes that the long-run effect of reduced incentives to do fundamental research could be a decline in the production of fundamental information.

Clark-Joseph (2013) uses an E-Mini Dataset to investigate mechanisms underlying the capacity of aggressive HFTs to profitably predict price movements. He notes that predicting the direction of future aggressive order flow in the E-Mini is relatively easy, but that the price impact of such predictable order imbalances is usually too small for indiscriminate trading ahead of the imbalances to be profitable. He tests an "exploratory trading" model that focuses on the ability of HFTs to submit small aggressive orders to detect information about the nature of passive liquidity available on the order book.

Specifically, the purpose of exploratory trading is to learn when order book liquidity is vulnerable and the price impact of future aggressive order flows is likely to be large. Using private information about effects on the order book that is generated by its own small aggressive orders, the HFT is able to trade ahead of predictable order imbalances with large aggressive orders at only those times when it is profitable to do so – that is, when the price impact is high. In sum, the exploratory trading model is based on identifying when order flow imbalances coincide with a vulnerable liquidity supply, and then trading aggressively with large orders at times when the price impact of future aggressive order imbalances is likely to be high. This description of an exploratory trading model functionally identifies a type of momentum ignition strategy.

Clark-Joseph (2013) tests the exploratory trading model by examining the trading profits of 8 A-HFTs that profit from their aggressive trading overall and significantly outperform non-HFTs. As noted in Section III.A.2 above, the 8 A-HFTs, which represent 24% of E-Mini volume with a 59.2% aggressiveness ratio, are much more profitable than the remaining 22 B-HFTs in his paper, which represent 21% of E-Mini volume with a 35.9% aggressiveness ratio.[32]

Notably, the 8 A-HFTs all lose money on their smallest aggressive orders (20 contracts or less). To test the predictors of the exploratory trading model, the author examines the extent to which information about the changes in the order book following small aggressive orders can explain the profits that various traders earn on subsequent aggressive orders. Consistent with the use of exploratory trading strategies by A-HFTs, the author finds that order book changes immediately following the small aggressive orders of A-HFTs provide significant additional explanatory power for A-HFTs' performance on their larger aggressive orders, but not for other traders' performance. He cautions, however, that the E-Mini findings may not be relevant to other markets and, as an example, suggests that the equities markets may not exhibit the predictability in demand that makes exploratory trading viable in the E-Mini market.

In addition, Clark-Joseph (2013) notes that aggressive HFT exploratory trading strategies are bolstered by two other advantages – low latency information and high frequency information. Low latency market information enables aggressive HFTs to rapidly obtain information about market response to their aggressive orders before the value of the information degrades with time. The high frequency information advantage, in turn, is

---

[32] Clark-Joseph (2013) notes that, throughout the E-Mini market, assorted linkages exist between various trading accounts, such as where a single firm trades with multiple accounts. Accordingly, the trading account divisions in the E-Mini Dataset do not necessarily reflect all of the trading accounts of a single proprietary trading firm. Some A-HFT accounts and B-HFT accounts, and some A-HFT accounts and non-HFT accounts (i.e., classified as non-HFT accounts pursuant to the methodology noted in Section III.A.2 above) belong to the same firms. As a result, various B-HFT accounts and non-HFT accounts may be either directly informed or able to make educated inferences about what one or more A-HFT accounts do.

inherent in the nature of HFT. Given that almost any aggressive order generates some amount of private exploratory information for the order submitter, a trader that places aggressive orders in greater numbers will gain access to greater amounts of exploratory information than less active traders that are not continuously engaged in the market. Moreover, as the number of orders increases, the time interval between a trader's aggressive orders also shrinks, so that the exploratory information produced by each order tends to be more valuable to the trader. The author states that these synergistic effects magnify the potential gains from exploratory trading for traders who place large numbers of aggressive orders.

4. Flash Crash – HFT Activity During a Severe Market Disruption

The events of the Flash Crash are described extensively in CFTC and SEC Staff (2010). These events will be discussed in a forthcoming part of this literature review that will address volatility more generally, including papers subsequent to CFTC and SEC Staff (2010) that shed further light on the dynamics of trading during the Flash Crash. This Part II of the literature review will focus on the three datasets of HFT activity during the Flash Crash that are included in Kirilenko et al. (2011) and CFTC and SEC Staff (2010).

To briefly summarize, the Flash Crash can best be described in terms of two liquidity crises – one at the broad index level in the E-Mini beginning about 2:41, and a second in individual equities beginning about 2:45.[33] During the E-Mini crisis, the E-Mini and broad market equity indices experienced rapid declines of approximately 5-6%, but then quickly reversed direction and recovered their declines. During the second crisis, prices in approximately 14% of individual equities, primarily ETPs, suddenly declined (increased) in an idiosyncratic fashion to as low as one penny (as high as $100,000), and then also rapidly reversed their declines (increases). By 3:00, both the broad market indices and individual securities generally were trading at price levels consistent with price levels prior to the two liquidity crises.

A prior version of Kirilenko et al. (2011) provided much of the basis for the discussion of HFT activity in the E-Mini in CFTC and SEC Staff (2010). They find that HFTs initially bought contracts from fundamental sellers in the E-Mini during the price decline, but then proceeded to sell contracts and compete for liquidity with fundamental sellers.[34] Based on this analysis, Kirilenko et al. (2011) conclude that HFTs did not trigger the Flash Crash, but that their responses to the unusually large selling pressure that day exacerbated market volatility.

As discussed in Section II.A.2 above, however, the analysis of HFT activity in Kirilenko et al. (2011) likely does not capture a large portion of total HFT activity in the E-Mini.

33 All time references are p.m. Eastern Time.

34 Kirilenko et al. (2011) defined fundamental buyers and sellers as trading accounts with an end of day net position no smaller than 15% of their trading volume on May 6, 2010.

Baron, Brogaard and Kirilenko (2012), for example, classified HFT traders as representing 54% of volume in the E-Mini, while Kirilenko et al. (2011) classified HFT traders as representing 34% of volume in the E-Mini. In particular, Kirilenko et al. (2011) adopted quite narrow metrics for defining HFT with respect to intraday and end-of-day net positions. As a result, the analysis of HFT activity in Kirilenko et al. (2011) (and therefore also in CFTC and SEC Staff (2010)) likely does not capture more than 1/3rd of HFT activity in the E-Mini during the Flash Crash, representing approximately 20% of total E-Mini volume.[35]

Turning to the equity markets, the FINRA Dataset in CFTC and SEC Staff (2010) reveals that 6 of the 12 HFTs in the dataset scaled back their trading. Notably, however, they did not do so until some point after the broad indices hit their lows at about 2:45. Two HFTs stopped trading at about 2:47 and remained inactive through the rest of the day. Four other HFTs appeared to have each significantly curtailed trading for shorter periods of time, ranging from as little as 1 minute to as long as 21 minutes. The scale back in HFT trading following the broad market decline appeared to have contributed to the evaporation of liquidity in individual equities. For example, the FINRA Dataset revealed that the HFT rate of participation in securities with broken trades was much lower than in securities without broken trades.

The Lit Venue Dataset in CFTC and SEC Staff (2010) confirms that HFTs in the securities markets did not curtail their trading until after the broad market price decline. The HFT percentage of double-counted volume increased from 42.6% during the 15 minute period ending at 2:00 to 50.3% during the 15 minute period of the broad market price decline ending at 2:45. During this price decline, HFT firms in the aggregate were aggressive net sellers of equities in the amount of $1.34 billion. In contrast, the HFT percentage of double-counted volume fell sharply to 36.6% during the next 15 minute period of the price recovery ending at 3:00. In sum, HFTs increased their activity, particularly their aggressive selling, during the period of rapid price decline in broad market prices, and curtailed their activity during the period when broad market prices recovered.

IV. Questions for Consideration

This is the second of a series to review recent economic literature on equity market structure. As noted above, the SEC staff review summarizes those economic papers that analyze recent financial data (2007 and later) and reach findings that in the staff's view are most relevant to important market structure issues facing the SEC. The staff's hope is that the literature review will help promote a dynamic exchange on market structure with and among the public, including investors, academics, and market participants.

With respect to this review, does it accurately describe the economic literature dealing with HFT, both with respect to its summary and discussion of the papers? Are there

35 The forthcoming Part III of this literature review will include papers that further examine the nature of trading in the E-Mini during the Flash Crash.

papers other than those identified in the review that should be included? In addition, most of the papers described above are working papers, which have not yet been through a peer review process and are subject to change as the authors respond to feedback. How reliable are all of the results of the papers discussed above? Are some results less reliable and why? Finally, what is the usefulness and applicability of the papers' analyses, metrics, and findings for the specific purpose of the staff's continued consideration of equity market structure issues?

V. List of References

A. NASDAQ Datasets

Brogaard,[36] Jonathan, Terrence Hendershott and Ryan Riordan, 2013, High frequency trading and price discovery, European Central Bank Working Paper Series No. 1602.

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

Gai, Jiading, Chen Yao and Mao Ye, 2013, The externalities of high-frequency trading, working paper.

Gao, Cheng and Bruce Mizrach, 2013, Market quality breakdowns in equities, working paper.

Gao, Cheng and Bruce Mizrach, 2011, High frequency trading in the equity markets during large scale asset purchases, working paper.

Gerig, Austin, 2012, High-frequency trading synchronizes prices in financial markets, working paper.

Hirschey, Nicholas H., 2013, Do high-frequency traders anticipate buying and selling pressure?, working paper.

Johnson, Hardy, Bonnie F. Van Ness and Robert A. Van Ness, 2013, Are all odd-lots the same? An analysis of all, pure, and circumstantial odd-lot transactions, working paper.

O'Hara, Maureen, Chen Yao and Mao Ye, 2012, What's not there: the odd-lost bias in TAQ Data, working paper.

Tong, Lin, 2013, A blessing or a curse? The impact of high frequency trading on institutional investors, working paper.

---

[36] This paper subsumes the results of Brogaard's earlier papers, "The Activity of High Frequency Traders", "High Frequency Trading and Market Quality," and "High Frequency Trading and Volatility." For further information, see the author's webpage at http://faculty.washington.edu/brogaard/research.shtml.

Zhang, S. Sarah, 2013, Need for speed: an empirical analysis of hard and soft information in a high-frequency world, working paper.

Zhang, Sarah and Ryan Riordan, 2011, Technology and market quality: the case of high frequency trading, ECIS 2011 Proceedings, Paper 95.

B. E-Mini and Flash Crash Datasets

Baron, Matthew, Jonathan Brogaard and Andrei Kirilenko, 2012, The trading profits of high frequency traders, working paper.

Clark-Joseph, Adam D., 2013, Exploratory trading, job market paper.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi and Tugkan Tuzun, 2011, The flash crash: the impact of high frequency trading on an electronic market, working paper.

Staffs of the U.S. Commodity Futures Trading Commission and the U.S. Securities and Exchange Commission, 2010, Findings regarding the market events of May 6, 2010.

C. International Datasets

Alampieski, Kiril and Andrew Lepone, 2012, High frequency trading in UK equity markets: evidence surrounding the US market open, working paper.

Alampieski, Kiril and Andrew Lepone, 2011, High frequency trading firms, order book participation and liquidity supply during periods of heightened adverse selection risk: evidence from LSE, BATS, and Chi-X, working paper.

Australian Securities & Investments Commission ("ASIC"), 2013, Report 331: Dark liquidity and high-frequency trading.

Benos, Evangelos and Satchit Sagade, 2013, High-frequency trading behaviour and its impact on market quality: evidence from the UK equity market, working paper.

Bershova, Nataliya and Dmitry Rakhlin, 2013, High-frequency trading and long-term investors: a view from the buy-side, *Journal of Investment Strategies*, 2(2), 25-69.

Breckenfelder, H.-Johannes, 2013, Competition between high-frequency traders, and market quality, job market paper.

Brogaard, Jonathan, Bjorn Hagstromer, Lars Norden and Ryan Riordan, Trading fast and slow: colocation and market quality, 2013, working paper.

Brogaard, Jonathan, Terrence Hendershott, Stefan Hunt, Torben Latza, Lucas Pedace and Carla Ysusi, 2012, High-frequency trading and the execution costs of institutional investors, Driver Review 21, Foresight Project on the Future of Computer Trading in Financial Markets.

Hagstromer, Bjorn and Lars Norden, 2013, The diversity of high-frequency traders, *Journal of Financial Markets*, 16(4), 741-770.

Hagstromer, Bjorn, Lars Norden and Dong Zhang, 2013, How aggressive are high-frequency traders?, working paper.

Investment Industry Regulatory Organization of Canada ("IIROC"), 2012, The HOT study: phases I and II of IIROC's study of high frequency trading activity on Canadian equity marketplaces.

Jovanovic, Boyan and Albert J. Menkveld, 2012, Middlemen in limit-order markets, working paper.

Kang, Jangkoo and Jeongwoo Shin, 2012, The role of high frequency traders in electronic limit order markets, working paper.

Malinova, Katya, Andreas Park and Ryan Riordan, 2013, Do retail traders suffer from high frequency traders?, working paper.

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