An Empirical detection of HFT strategies [preliminary]

October 1, 2015 Abstract

This paper detects empirically the presence of High Frequency Trading strategies from public data and examines their impact on financial markets. The objective is to provide a structured and strategic approach to isolate signal from noise in a high frequency setting. In order to prove the suitability of the proposed approach, several HFT strategies are evaluated on the basis of their market impact, performance and main characteristics.

1 Introduction

High-frequency trading (HFT) refers to a subset of algorithmic trading that involves the usage of specialized digital infrastructure and ultra-low latency computing to execute a wide range of strategies on financial markets (OHara 2015). While HFT has been present since the digitization of markets, it has only become wide spread over the past decade. Evidence suggests that a very large percentage of the trading activity in many markets is now carried out via HFT algorithms across a variety of assets. Projections indicate that this presence is likely here to stay, and perhaps even to increase over the coming years (Kumar, Goldstein, Graves, 2011). HFT has been the recent focus of researchers, particularly those focusing on market microstructure, as well as regulators and even the popular media, mainly owing to increasingly frequent incidents of abnormal market behavior supposedly linked to HFT algorithms. Examples of such events include the Flash Crash of 2010, which has since attracted investigations by multiple researchers, media reports and has even inspired entire books (Lewis 2014). Similar, but less high profile, accidents have also occurred in other markets, namely the treasury markets in October 2015 and on the fixed income instruments. The impact has not been homogenous, and in some instances, prices have actually increased abnormally. This was the case with a recent episode with US Treasuries, where the market saw a flash crash of yields indicating a rapid increase in price as discussed in the U.S. Department of Treasury, 2015. Whether these accidents are entirely to blame on HFT activity is not yet certain, but it has captured the interests of regulators, investors and market participants. Hence, a deeper understanding of their working and impact on the wider markets and asset prices is crucial. The HFT paradigm shifts for markets makes necessary the usage of new tools to detect and analyze the activities of high speed traders. Such tools may be of interest to regulators, investors and academics alike. The methods presented in the current paper provide a detailed strategic approach which aims at isolating signal from noise. Appropriate benchmarks are selected to evaluate the performance and the impact of two different HFT strategies.

Weighing up the costs and benefits to financial markets arising from HFT activities is an exceptionally complex task. On the one hand, lightning fast trading may mean that markets are becoming increasingly informationally efficient, as news and data are rapidly incorporated into prices. Efficient capital markets are of crucial importance to the optimal allocation of scarce resources in an economy. By utilizing the pricing mechanism, efficient and a well functioning capital markets bridge the gap between economic agents with a deficit of funds and those with a surplus and ensure that wealth is employed where it produces maximal returns. On the other hand, critics of HFT often argue that they contribute to the noise in the market, and that their activities are disconnected from economic or financial fundamentals, creating volatility which discourages long term investors and issuers of securities. Others yet argue that HFTs are simply deliberately set up to game the system and exploit opportunities arising from market microstructure as organized by exchanges in order to profit at the expense of other participants. Many of these arguments appear to have some merit (Turbeville 2013).

The debate between supporters and opponents of HFT is heated and often confusing. The powerful interests of lobbyists and vested stakeholders on both sides (HFT firms on one end, buy-side investors and traditional traders on the other) are in constant battle to persuade regulators, exchanges and the public to their views (Rijper, Sprenkeler, and Kip 2010). Joining this debate is not the objective of the present paper. Instead, it deliberately focuses on providing a positive and empirically driven point of view.

2 HFT strategies and their impact

One of the main challenges for researchers on the topic of HFT is the wide diversity of strategies employed by practitioners these may range from traditional market making through statistical arbitrage to more speculative and aggressive strategies.

Traditional market making is well researched and involves a market participant taking both sides in the market for a security simultaneously, by providing pairs of Bids and Offers. A market maker aims to make a profit based on the spread between the Bid and Ask prices, and in the process needs to manage inventory and execution so as to minimize undesired exposure to price risk. Market makers have been present on markets for a long time, and in fact in some cases have been institutionalized such as historically the case of the NYSE specialists, the Designated Market Makers (DMMs) of Euronext and the Kursmakler of Frankfurt Stock Exchange. In such arrangements, the market makers are allowed to profit from an exclusive proprietary knowledge of the security and the order flow relating to it. In exchange, they are often expected to maintain an orderly market in the security a duty to step in and stabilize the market in cases of extreme price fluctuations, perhaps even at the costs of incurring a loss in the process. Traditional market microstructure models have been utilized to understand the behavior of market makers and the factors influencing their decisions see for instance, (Kyle 1985).

HFT practitioners which chose to act as market makers in securities may follow the exact same strategies as their traditional counterparts with several significant differences. The main distinction of HFT market makers is their mandate for opportunistic behavior and their lack of long term obligations to market participants. While traditional market makers are well known and have certain responsibilities for the markets they make, HFT firms are often small and relatively unknown entities. They may withdraw from the market at any point if they feel that conditions are unfavorable. In fact, such decisions may be taken automatically by the algorithms which can detect changes in market conditions much more rapidly than their human counterparts. Indeed, this is what many argue may have occurred during the Flash Crash where a deterioration in market conditions may have triggered algorithms to withdraw from the market, causing further deterioration in liquidity, in a vicious cycle. On the other hand, prices quoted by HFT market makers may be quicker to adjust to news regarding the fundamentals of the underlying security, or the information conveyed by rapid changes in order flow. Similarly, from a theoretical point of view HFT may be better at managing inventories in real time, thus reducing the inventory risk component priced in quoted spreads. Indeed, it is common for HFT firms to carry flat positions overnight regardless of the strategy they follow. Another benefit for market efficiency should be the instantaneous incorporation of cross asset arbitrage into pricing in real time which may only be achievable through the usage of low latency algorithms.

Statistical arbitrage, pattern recognition and directional trading are also well known strategies which have been employed widely before the rise to prominence of HFT. These are mainly aimed at exploiting long lasting patterns and relationships in asset prices in order to profit. Such strategies are often employed by aggressive buy and sell-side institutions such as hedge funds or the quantitative strategy desks of investment banks. While many of these techniques may be carried out through traditional algorithmic trading, it is clear that HFT may provide some incremental benefits, due to the speed of execution and detection of the patterns, which is always valuable in highly competitive financial markets.

The last group of strategies employed by HFTs are exclusive to them and thus form a very distinct group which is separate from the previously examined ones. This group of strategies have been called predatory, malicious and manipulative by researchers and market participants and may be seen as exploiting market microstructure at the expense of other traders, see (Foucault, Hombert, and Roşu 2015) and (Biais, Foucault, and Moinas 2015) and (Biais and Foucault 2014).

Some of the most commonly referenced predatory strategies include Quote Stuffing, Momentum Ignition, Order Fade and Pinging. Quote Stuffing is a strategy characterized by HF traders rapidly submitting a torrent of orders (bid, ask or both simultaneously) in the order book over an extremely short period of time. Reports by Credit Suisse estimate that the majority of Quote Stuffing episodes last up to 2 seconds. However the reported distribution is uneven, with some episodes lasting significantly longer. These events are frequently observed in various markets and assets. An obvious objective is to flood the market systems with (useless) information in the form of rapidly changing quotes, so as to slow down and confuse the responses of competing algorithms and other traders. This is very similar to a DDoS attack on a computer system. However, there are other possible objectives. One measure frequently used by algorithmic traders as well as traditional investors in benchmarking trade execution and market prices is the midpoint this is derived as the average of the Bid and Ask prices and is then used to determine the relative spread or the ex-post quality of trade execution. By practicing Quote Stuffing, HFT may induce significant quote volatility which may cause the midpoint to also fluctuate. These dynamics may induce certain algorithmic or human traders to trade based on the rapidly moving midpoint values. However, in the instance that the trader carelessly utilizes market orders instead of limit orders, they may receive execution at a different price than anticipated. During Quote Stuffing episodes for instance, the bid and/or ask quotes rapidly fluctuate over the space of microseconds. In the case where the underlying asset is also traded in Dark Pools which use the midpoint in lit markets as part of their derivative pricing mechanism, HFT may attempt to manipulate the midpoint via Quote Stuffing in order to gain favorable execution in the Dark Pool.

Quote stuffing is easily observed via several characteristic patterns of quote volatility which may occur on the Ask, the Bid or both simultaneously. Empirically, Quote Stuffing is observed as having some influence on the direction of price moves immediately following an episode with prices seen as more likely to move in the direction of the stuffing. Therefore, when the stuffing occurs on the Bid side, prices are observed as they are more likely to fall; whereas the opposite holds if stuffing occurs on the Ask side. Quote Stuffing is a very frequently observed phenomenon, up to 53 occurrences per day in some market segments as reported by (Tse, Lin, and Vincent 2012).

Momentum Ignition is a strategy which is used by HF traders to cause and exploit sharp and rapid movements in prices occurring over a limited amount of time. (Tse, Lin, and Vincent 2012) observe that the average Momentum Ignition events occur over the space of 1.5 minutes. The aim of this strategy is to instigate other traders to join the speculative movement and cause a rapid snowball like price cascade. This is achieved through trading in high volume in specific patterns aimed at triggering other market participants to follow. If successful, the strategy results in a sharp price move which is usually not of very large size; although it is significant when the short time frame of the episode is taken into account. Once prices have moved, the HFT will cover their position and the move will reverse itself albeit on much lower volume than the original. A fully executed episode of Momentum Ignition leaves a distinct impact on price and volume data. This is characterized by a spike in volume accompanied by flat prices. This is then followed by a sharp price jointly with an even greater increase in volume. Once the high or low point of the price move is achieved, volume declines and price gradually reverts back to its original level. These events are relatively less frequently observed with an average of at least 1.6 event per stock per day as reported by (Tse, Lin, and Vincent 2012). This pattern is of particular interest for the current paper as it carries a direct and readily measurable economic result for practitioners, while assuming seemingly very low risk as the move is caused and exploited at the same time by HFT.

Order Fade and Layering refers to two very similar strategies used by HF traders to front run large order volume or create a false illusion of liquidity within the order book. Practitioners will cancel partially (for some of the volume) or completely their orders on different levels of the order book, following the execution of a trade. The aim is to make the traders demanding liquidity and thus pay a higher price by having to run up or down the book. This is also, however, an example of fleeting liquidity caused by HFT. The orders submitted in the order book disappear before they can be executed, therefore creating an illusion of liquidity which is not really there (since it cant be transacted on). Similarly, when Layering, the HF traders will submit a large number of orders on the bid or ask side of the order book in order to create a false impression of strong buying or selling pressure and manipulate other traders. The instigators however have no real intention of executing these orders which will be cancelled by the low latency algorithms as soon as execution is attempted. Both of these practices are extremely frequently observed and have become widespread in many markets. Observing and analyzing these strategies fully requires the usage of more than one level of the order book and is beyond the scope of the current research. Further factors which complicate the analysis are data synchronicity issues and cross venue fade, where events on one trading venue cause the algorithms to adjust their position in the order book on a different venue simultaneously.

Pinging refers to a practice where the high speed traders place a series of small volume orders with the objective to survey the market and hunt for large institutional traders who are seeking to trade for liquidity reasons. This is done by sending orders within the frames of the prevailing bid-ask spread at the time in order to entice the large uninformed traders to react and execute. Once this is detected, the algorithms deplete all liquidity on the side of the order book the large trader is interested in transacting with. If the large trader is looking to sell for instance, the algorithm will look to dry up all the bid liquidity in the market over a very short period of time and replace it with its own orders at a much higher spread than before. This is a tactic used to force large traders to pay a higher spread, and once executed, the algorithms will unload their position at the much more favorable market prevailing spreads. This tactic of front running is one of the most frequently cited by critics of HFT and is referred to as whale hunting, as in (Turbeville 2013).

Clearly, in a complex and fragmented market place there can be no real end to the strategies employed by market participants, and even more so for highly sophisticated ones such as algorithmic and HF traders. For the sake of maintaining the focus of the present research paper, only two of these strategies, Quote Stuffing and Momentum Ignition are analyzed in detail. Hopefully, the techniques and analytical tools provided in this paper may be adapted and used by future researchers to analyze a further array of strategies encountered in the market place.

3 Selection of assets and data sample

3.1 Selection of assets

The main contribution of this paper lies on its strong empirical focus to offer an array of ready tools aimed at detecting and examining the presence of HFT strategies from the public market data. The main immediate prerequisite for examining HFT strategies is the access to high quality data. As there is ample evidence of the presence of high frequency traders in the market across various markets and assets, we select a limited but diverge range of assets for testing the new measures. Upon their successful use, the data sample may be readily expanded for future research to include additional assets. The assets selected are: Apple Inc. (ticker AAPL), Bund Futures contracts maturing in September 2015, and the US Oil ETF (ticker USO). This is a diverse range of assets, encompassing an equity instrument, a futures contract on German 10 year government bonds (fixed income), and an ETF closely related to a major commodity. Furthermore, each of these assets is often anecdotally hypothesized by market participants to be favorite for HF traders. For instance, it is known that during the events of the 2010 Flash Crash, Apple stock was briefly driven up in value to as high as \$ 100 000 within a few instantaneous trades by malfunctioning algorithms, while the majority of the other assets were collapsing (Patterson 2013). Apple is also the most actively traded share in the world, with an average daily volume of over 63 million shares in the last 50 days. It means that Apple is likely to attract high levels of activity from a large amount of diverse market participants, including the HF traders. Furthermore, the stock is included as a component of several other actively products with a heavy weighting, such as QQQ ETF focused on the technology sector, and the Nasdaq index. This

suggests that cross asset arbitrage which is the favorite HFT strategy is also available. Additionally, being a tech-stock, the stock exhibits sufficient volatility to make short term price fluctuations plausible and thus making it potentially more attractive to HFT activities such as Momentum Ignition. The period of study selected for Apple is the week spanning from January 26th till 30th of 2015 around the earning report release. A week marked by a scheduled significant catalyst event is likely to be a week of increased trading activity by investors. The large volume may then draw the attention of HF traders interested in employing predatory strategies.

Similarly, traditional market makers and proprietary trading firms often lament the increasing activity of HFT in the market for Bund contracts, and express their difficulties of competing against their high speed algorithms. The Bund futures contract which is traded on Eurex is a contract based on Fixed Income Securities issued by the German Government. This can be for the delivery of a notional debt obligation of the Federal Republic of Germany with a time to maturity between 8.5 and 10.5 years, and a notional coupon rate of 6%. The contracts are issued in March, June, September and December. This asset is extremely popular with traditional proprietary trading firms and market makers, and it is considered one of the most accurate indicators of the prevailing interest rates in the Eurozone. As such, it is heavily influenced by the monetary policy decision of the European Central Bank, as well as very sensitive to data about the variety of macroeconomic indicators such as inflation, economic growth, unemployment, sovereign debt levels. The week selected for this study spans from June 1st till June 5th of 2015, as this week has been marked by significantly high volatility in European Fixed Income markets, referred to "bloodath". During this week, the monthly monetary policy decisions and press conference were hosted by the ECB on the 3rd of June. This particular day saws a continuation of the sharp moves in the Fixed Income markets raising the question whether sovereign yields volatility being raised during the press conference. The ECB president highlighed technical factors, including the volatility clustering and the poor liquidity conditions as some of the possible causes of the moves. Nevertheless, the abnormal price reaction of the week raises the question of possible HFT involvement, perhaps being drawn by the extreme volatility or even instigating it. Therefore, a week of heightened trading activity and the presence of a strong catalyst event vis-à-vis the monetary policy announcements may be a good starting place to search for the footprint of HFT strategies.

Finally, the use of US Oil ETF (ticker USO) is particularly relevant, as it is reported as one of the top holdings of major HFT firms ¹, such as the Virtu Financial. The US Oil

 $[\]label{eq:linear} {}^{1}\mbox{http://www.bloomberg.com/news/articles/2015-02-19/berkshire-hathaway-exotic-etfs-among-flash-boy-holdings}.$

ETF is an ETF designed to track the daily movements of WTI light, sweet crude oil. The objective of the ETF is for the daily changes in its net asset value (NAV) to reflect the daily relative changes in the price of WTI as measured by a Benchmark Futures Contract traded on NYMEX. The product trades in \$0.01 increments, and incurs a 0.45% management fee by its administration company, Brown Brothers Harriman Co. Much more significantly, however, USO is listed as the second largest holding for prominent HFT firm Virtu Financial as of the beginning of 2015. This has several important implications in the present context. Traditionally, the aim of Market Making Firms like Virtu is to manage and minimize inventory risk, particularly arising from overnight exposure. In order to achieve this, it is unusual for them to hold overnight positions in securities. This only makes sense in a case where a particular security is very heavily focused on, and therefore warrants incurring some overnight exposure in order to facilitate smoothness market making operations. This is done to ensure that an appropriate level of inventory is maintained at all times. Furthermore, anecdotal evidence from the industry suggests that ETF and Delta One product arbitrage is one of the most popular HFT strategies for many firms, including the likes of Flow Traders, Optiver and Virtu. The basic premise of this family of strategies is to exploit transient discrepancies between the NAV of an ETF product, and the value of its underlying basket of benchmark assets. This type of cross-asset arbitrage is a possible strategy with assets such as USO. Finally, as with the other assets examined in the present research, there are significant fundamental catalyst events occurring for the security during the time frame used. While oil prices have experienced elevated volatility since the beginning of 2014, July 2015 marks an important geopolitical development for Oil markets. After several days of intense high-profile negotiations between the Iran Government and the international community, a deal was reached to lift sanctions on Iranian Oil exports, in exchange for restrictions in the countrys nuclear programme. Fundamentally, this development has the potential to cause significant volatility in oil markets due to changes in supply to western markets. It can be argued that events of this magnitude may have enticed strategic portfolio re-balancing by commodity traders. This could also mean an increase in the activity of HFT market makers. Therefore, for the purposes of the present research, data is used on USO trades and Level 1 quotes for the 13th and the 14th of July 2015.

3.2 Data

The data for the 3 examined assets in the project shares the same basic properties. It consists of Trade and Level 1 Quote data. For Trades, data includes: Trade price, Trade quantity and Trade Time. For Level 1 Quotes, it includes Bid Price, Ask Price, Bid Quantity, Ask Quantity and Quote Update Time. The data is consolidated from multiple venues and spans across pre-market trading and after Hours, as well as regular trading sessions. The precision of the time measurement is at the microsecond levels, whereas for prices it is determined by the relevant minimum price increment. All of this ensures the high accuracy of the information examined, and minimizes the risks of problems such as loss of synchronization manifesting. The order to trade ratio across the week investigated is on average of 11%.

A central challenge which emerges when working with High frequency data over a large time period is processing the vast amount of information. The task of finding short episodes of HFT activity in the enormous amounts of data used is similar to the proverbial needle in a haystack. This makes clear the need for a deliberate strategy of analysis aimed at isolating signal from noise. Otherwise, the only alternative would be a manual approach whereas the researcher examines personally the data. A quick simple calculation, assuming that trades and quote updates are examined in batches of 10 consecutive and that each batch is processed by the researcher in just 1 second (speed which is most likely unachievable in reality) indicates that it would take a total of over 266 hours to process the entire sample used in the current project. More realistic assumptions could result in a processing time which is even 30-60x times longer. Clearly, this approach is not viable, and a better, more efficient solution is needed. An obvious starting point is to look at the algorithmic techniques used by high speed traders themselves and to utilize similar tools for the purposes of the research. This is the main contribution of the present paper.

To begin processing the data, it is necessary to reflect on several important properties of the instances of HFT activity. Namely, these are usually seen as occurring over a specific time interval marked by a starting point, a time span, and an ending point. Similarly, as these are driven by algorithms sensitive to market conditions, the period immediately preceding an outburst of HFT activity might also be of particular interest in the analysis. The time spans over which HFT events last can be quite flexible and may follow a fat tailed distribution as indicated by (Tse, Lin, and Vincent 2012). This means that a one-size fits all approach could be wrong, and a certain level of flexibility is warranted. Additionally, different types of events might occur over drastically different time horizons, for instance while Quote Stuffing may only last several seconds in most cases, and Momentum Ignition typically spans over several minutes. A truly robust strategy for HFT detection would necessitate a sufficient built in scalability to cope with this without any fundamental alterations. It is also important to keep track of market and institutional conventions, such as the difference between Pre-market, regular hours and After Hours trading, which will have an obvious and profound impact on the level of activity during different times of the day. All of these considerations need to be built into the analytical strategy to ensure that it is appropriate for the current analysis.

4 Detection of HFT strategies

The raw data is received in plain text format and includes multiple variables relating to trades and quotes. This is processed by the analytical tool in order to extract only the entries relevant for the purposes of the present analysis. These are stored as variables. As a consequence, it is immediately possible to extract statistics. For the purposes of analyzing HFT events, the best approach is to use the method of rolling time-frames with overlap. Since the data points are unevenly distributed in time, an algorithm is used to collate them into subsamples (windows) spanning a user specified time length. Therefore, each data point serves as a starting point for a window which includes a number of other data points which fall within a pre-specified time from the first one. This is a flexible approach which allows the data to be analyzed frame by frame, and for statistics to be estimated for each window. Furthermore, since the windows are derived with overlap, this means that all potential frames are examined. A big advantage of this approach is that the researcher is able to specify a desired window size (in microseconds), which can mean that HFT events of differing length can be analyzed by the tool. The window subsample framework is at the center of all the detection algorithms used in this paper.

Detection of HFT activity in large data samples is the main objective of the current research. As noted previously, this is a challenging task which calls for an algorithmic solution where little or no human input is required, simply owing to the volumes of data examined. Fortunately, the software tool used in this project provides an efficient environment that can be used to implement such algorithms. The time window framework allows for statistics to be estimated for each of the rolling subsamples. This simplifies the task of detecting the time intervals containing HFT activity to designing statistics which capture the similarity of the pattern of activity observed in a given subsample to that of typical HFT patterns. The one arbitrary element in this strategy is the length of time frames for both Quote Stuffing and Momentum Ignition. Alternatively, these can be determined empirically by trial and error.

4.1 Quote Stuffing

Quote stuffing is an event during which high quote volatility is observed over a very short period of time, typically lasting several seconds. During this time, rapid but transient quote updates occur, often following several specific patterns. Figures 1 and 2 show some of these patterns.



Figure 3: Bund Combined stuffing

Additionally, Quote Stuffing can occur either on the bid side, the ask side of the market or both simultaneously as shown in figures 1 to 3 above. In order to detect the presence of these patterns of activity in the subsample contained in the data, the D-ratio statistic is estimated. This can take on several alternative specifications depending on the type of Quote Stuffing observed bid, ask or combined. The D-ratio is a geometry based metric, that is inspired by the graphical representation of Quote Stuffing on a chart. In essence, all Quote Stuffing episodes share several key characteristics, irrespective of the particular pattern.

- They include small rapid movements in the bid, ask or both levels, which are subsequently rapidly reversed.
- This is repeated a large number of times, over a small time frame.
- Over the span of the entire time frame, the actual direction movement in the quote levels is low, if any.

The D-ratio has four components: (i) Carryask; (ii) Bigask; (iii) Carrybid; (iv) Bigbid. The window (subsample) size is denoted as j, each Bid quote price is denoted as B, and each Ask quote price is denoted as A. Therefore the four component variables are: The D-ratio has four components: (i) Carryask; (ii) Bigask; (iii) Carrybid; (iv) Bigbid. The window (subsample) size is denoted as j, each Bid quote price is denoted as B, and each Ask quote price is denoted as A. Therefore the four component variables are:

1. Carryask is the sum of absolute incremental (instant-by-instant) changes in the ask price over the period.

$$Carryask = \sum_{i=2}^{j} |Ai - A_{i-1}| \tag{1}$$

2. Bigask is the absolute change in the ask price level between the starting and the ending points of the period examined.

$$Bigask = |A_i - A_1| \tag{2}$$

3. Carrybid is the sum of absolute incremental (instant-by-instant) changes in the bid price over the period.

$$Carrybid = \sum_{i=2}^{j} |Bi - B_{i-1}| \tag{3}$$

4. Bigbid is the absolute change in the bid price level between the ending and the starting point of the period.

$$Bigbid = |B_j - B_1| \tag{4}$$

These above variables are used to compute three alternative specifications of the D-ratio:

• Ask specification aims at detecting the in-ask Quoting Stuffing activity. This implies rapid quote volatility on the ask side, and a relatively passive bid side. The specification is:

$$D = \frac{\frac{carryask}{bigask}}{\frac{carrybid}{bigbid}}$$

In order to guarantee the function's solutions domain, several special cases are defined: (i) if bigask=0, it is instead set at the level of the minimum tick increment at 0.01; (ii) if bigbid=0, it is instead set at the level of the minimum tick increment at 0.01; (iii) i carrybid=0, the entire denominator $\frac{carrybid}{bigbid}$ is set to equal to 0.01. This ensures that a corresponding D-ratio can be calculated for any given subsample.

• Bid specification aims at detecting Quote Stuffing which occurs on the bid side of the market. This is characterized by rapid quote volatility in the bid price levels, while the ask price remains relatively inactive. The specification is given by:

$$D = \frac{\frac{carrybid}{bigbid}}{\frac{carryask}{bigask}}$$

For solution domain reasons, several special cases are defined: (i) if bigask=0, this is instead set at the minimum tick increment at 0.01; if bigbid=0, this is instead set at the minimum tick increment at 0.01; if carryask=0, the entire denominator $\frac{carryask}{bigask}$ is set to 0.01 instead.

• Combined specification aims at detecting Quote Stuffing Activity of the combined type which occurs on both the ask and bid sides of the market. This is characterized by a period of high quote price volatility which occurs over a short period of time, but it is driven by transient movements. The specification is given by:

$$D = \frac{carryask}{bigask} + \frac{carrybid}{bigbid}$$

For the purpose of solutions domain considerations, several specific cases are predefined: If bigask = 0, it is instead set at the minimum incremental tick size at 0.01. If bigbid = 0, it is instead set at the minimum incremental tick size at 0.01. Figures below present some of examples of the results derived using the D-ratio.





Figure 6: Apple in-ask stuffing

The final step of the detection strategy, after the D-ratio values for each window in the sample have been calculated, is to determine which ones are indicating a potential period of HFT activity. Since a unique window (subsample) is associated with each data point in the sample, and a D-ratio value is associated with each window, it is possible to use the observed distribution of D-ratio values over the entire sample, and subsequently select a cut-off point for the most promising ones. The Trident tool supports a user specified cutoff point. Once a D-value has been estimated for each window in the sample, the entire array of D-ratio values is ordered in increasing order. The user specified percentage is then applied to select the cutoff point. This is done via the below formula:

$$D_{cutoff} = D_{arraysize-(rounddown(percentile*arraysize))}$$
(8)

The cutoff determined using this technique has the major benefit of coming from the distribution observed within the actual data, rather than an arbitrary level selected by the researcher. Since a higher D-ratio should indicate higher likelihood of HFT activity, using the observed distribution of the values within the sample allows the researcher to focus on a subsample that appears to be most relevant within a given context. Once the cutoff is determined, it is used to filter out only the windows (subsamples) which have a D-ratio value above the cutoff point. These can then be examined using additional automated techniques or manually by the researcher. For the purposes of most of the present analysis, the second approach is taken.

4.2 Momentum Ignition

Momentum Ignition is an HFT strategy that is characterized by a specific pattern observed in both trade prices and trade volumes. This usually comprises three main stages:

• An initial spike in trade volume, which isnt accompanied by any significant changes in price.

- A subsequent sharp price move (positive or negative), accompanied by a new, even larger increase in volume
- A gradual price reversal to levels observed before the event, accompanied by low volume.

This pattern is observed over significantly longer time frames than Quote Stuffing, and may last for up to several minutes. It is also relatively less frequent in occurrence, although still prevalent in most traded instruments at least once per day with higher activity in certain sub sectors of the market. The duration of the events, as well as the severity of their market impact appears to follow a fat-tailed distribution, with a small fraction of events having major market impact and lasting for a prolonged period of time. Unlike Quote Stuffing where most of the economic impact is indirect resulting from changes in the midpoint or enticing potential traders, Momentum Ignition has a directly observable economic impact, which can be measured in relative terms (size of the price move in bp), or, potentially even in absolute (price change x estimated position by HFT).

The strategy used for Quote Stuffing detection is modified and applied for detecting Momentum Ignition. While the characteristic pattern of Momentum Ignition includes two dimensions, trade prices as well as volume for the sake of computation efficiency and simplicity, only trade prices are used an input for the D-ratio specification used to detect the pattern. The software has additional built in capabilities which allow the researcher to easily examine the volume data for the characteristic patterns following the initial screening. Furthermore, it is assumed that, as consistent with previous empirical studies of financial markets, the distribution of asset returns exhibits leptokurtosis, and therefore a small enough fraction of subsamples (windows) will contain the large moves relevant for studying Momentum Ignition. The distribution derived strategy and the specification of the D-ratio used ensure that the biggest relevant price moves present in the data are examined. The D-ratio used for Momentum Ignition detection is based on 3 key inputs:

- 1. Starting trade price is the Trade price in the starting point of the time period.
- 2. Ending trade price is the Trade price of the final trade in the time period.
- 3. Price Span defined as |EndPrice StartPrice|. If this turns out to be 0, then it is set to 0.01 instead for domain purposes.

As with Quote Stuffing, the metric estimated is inspired by the geometry of a graphical representation of the pattern. In the case of Momentum Ignition, this involves estimating two distances for each trade (t) included in the subsample (window):

$$D1_t = |P_t - StartPrice|$$

$$D2_t = |P_t - EndPrice|$$
(9)

These metrics are used to derive a Total Distance:

$$TD1_t = D1_t + D2_t \tag{10}$$

Once this is derived for each trade in the subsample, the largest TD is determined:

$$TD_{max} = \max \sum_{i=0}^{n} TD_i \tag{11}$$

These inputs can be used to derive the value of the D-ratio for the window:

$$D_{sample} = \frac{TD_{max}}{PriceSpan} \tag{12}$$

Once a D-ratio value is estimated for each window in the sample, the array of D-ratio values is ordered and a cutoff point is determined. This is then used to filter out the top values encountered in the sample. The focus on price data only means that this approach will select the windows with the biggest price moves which have subsequently reversed back to their starting point. Once these are determined, the researcher can manually examine each of these and use the built in functionality of the Trident Tool to look for the characteristic pattern in volume, finally yielding a confirmed finding.

5 Methodology

While the detection of HFT activity is the main objective of the present research, the study is broadened by using further analytical tools to obtain additional information regarding the impact and performance of high speed traders on the market. The analytical suite built into the Trident Tool provides an array of ways to extract that information efficiently from the data. This includes metrics such as the Quoted Spread, the Level 1 Quoted Depth, Midpoint data, Trading Volume, Trade Volume Moving Average, VWAP and Implementation Shortfall.

• Quoted Spread is one of the most widely used indicators of trading costs in financial markets. While it is not necessarily indicative for very large orders, which may deplete Level 1 liquidity and walk up or down the book, it is a generally good indication for average trades.

- Level 1 Quoted depth is custom built indicator which captures the total quantity offered at the best ask and the total quantity demanded at the best bid. It is a good indicator of market liquidity as it shows the actual potential trade size which the market can transact without too much impact. It also appears to act as an indicator during quote stuffing episodes.
- The Midpoint is a popular indicator used by market researchers and practitioners alike. This is often seen as a metric indicating the relative performance of a trade, or could be used to derive more sophisticated measures of the spread. Furthermore it is incorporated into the trading strategies of institutional investors, and thus may influence the behavior of block trades. Finally a very important usage of the Midpoint is observed on some Dark Pools, which use a crossing rule based on the midpoint of the lit markets. These factors make the Midpoint one of the most important metrics for the practitioners of strategies such as Quote Stuffing.
- VWAP The Volume-Weighted Average Price is another extremely popular metric used by industry and academia alike. It is a measure derived for a specific time period, and may be used for the purpose of finding a realistic assessment of the most likely trade price over that time. It is often used as a simple benchmarking tool for the performance of trading desks at institutional investors. It is also used a benchmark in many algorithmic strategies. For instance a trader who manages to buy securities at below VWAP, and sell them at a price above VWAP during a certain time period, can be assumed as having contributed value. A trader who has followed VWAP without deviations is known as having delivered accurate execution. A popular alternative strategy, used by many institutional investors is the constant participation rate strategy, where a large order is worked by being split into child orders of small sized and executed throughout the day in proportion to the average trade sizes observed throughout a day. The VWAP can be seen as providing a more accurate view of the prices of securities, as it assigns a lesser weight to abnormal volatility stemming from low volume trades. In the context of the present analysis, the VWAP over a time period (window) is calculated recursively. This applied to recursively calculate VWAP at each moment within the window examined.
- Implementation Shortfall is a metric which builds onto VWAP, and is also widely used by market practitioners. This is aimed at measuring the execution performance of traders and algorithmic strategies, by benchmarking it against a hypothetical paper portfolio executed at the midpoint once the order is received. In the context

of the present research, this is coupled with VWAP to obtain of measure of market movement. The total quantity traded during the window examined is assumed as the order being worked by the market. The result is a metric following the price movements during the period, but is adjusted for the initial midpoint. It is calculated assuming a buyer point of view, therefore a positive value indicates that a buyer would have been better off executing immediately at the midpoint at the start of the time period examined (window), rather than delay execution partially or fully. Similarly, negative values indicate that the price moves lower through the window so from a seller point of view, it is ideal to execute immediately.

- Trading volume is displayed as a metric which is crucial for the correct identification of Momentum Ignition patterns. This allows the researcher to quickly identify trends in volume.
- Trade Volume Moving Average is a built in metric which is aimed at simplifying detecting trends within large amounts of trade volume data. This is particularly useful when examining time windows containing a large number of data points. The moving sample for the average is set at 100.



Figure 9: Bund Level 1 Quoted Depth



Figure 15: Apple VWAP



Figure 17: Apple Implementation Shortfall

5.1 Artificial Neural Network

09-26-11 677

The final objective in the present research is an attempt at detecting potential commonalities which may signal that an episode of HFT activity may be imminent. This is achieved through performing an Artificial Neural Networks experiment. There may be reasons to believe that at least part of the HFT activity may be predictable to some extent. Fundamentally, algorithmic trading relies on an algorithms being triggered by market conditions. If these conditions were known, it would also be possible to forecast when algorithmic activity may be imminent. However, in reality, this information would constitute a very closely guarded company secret, and is almost certain to be protected as intellectual property. Therefore, an alternative method is to attempt to estimate and detect commonalities in market conditions immediately preceding an episode of HFT activity.

Artificial Neural Networks (ANN) are types of statistical learning models which are designed in a way that mimics the logical structure of a biological brain. These models are particularly good for the purposes of pattern recognition, and are used in a variety of applications, including speech recognition, image recognition, analyzing patterns of consumer behavior and in financial markets. Fundamentally, all ANN models require at least two basic characteristics a topology and a transfer function (Bishop 1995). ANNs are constructed out of nodes called neurons which act as simple I-O transformers. Data is fed into neurons as a signal input, and this is processed via a transfer function which generates an output signal. There are multiple transfer functions available, which have different characteristics and may be appropriate for analyzing specific problems. Some of the most widely used ones include the logistic function, the linear function, a hyperbolic function and a threshold function. For instance, a logistic transfer function implies that the value of the potential outputs may range between 0 and 1. The derivative of the transfer function has important implications for the performance of the module during learning on training data sets. For the purposes of the present research, the hyperbolic function is used:

$$O = tanh(I) \tag{13}$$

The derivative of the hyperbolic function is approximated by:

$$1 - x^2 \tag{14}$$

This ensures that outputs can take on values between 1 and -1, as can be see below. Additionally, a large central region of the function is characterized by a relatively constant slope, allowing for strong learning performance in a wider region of input values:



Figure 18: Plot of the hyperbolic transfer function

The neurons of an ANN are structured in groups called layers, and while there is no limitation to the number of layers, and the number of neurons in a layer, most topologies will consist of 3 layers an input layer, a hidden layer, and output layer. Each neuron in a layer is connected to all the neurons on the layer immediately preceding it, and to an additional bias neuron, which has a constant output. These connections are assigned to a specific weight each, and the weighted sum of the signals coming from all connections forms the total input. The input layer neurons are used as input nodes, where raw data feeds into the network directly. This is then processed via the transfer function of the neurons, and fed via connections to the hidden layer, which then processes the signal and transmits it to the output layer. The output layer generates the final output of the network. Training is a key stage of using an ANN in a practical solution. One of the major advantages of using this type of model, is that it doesn't require prior knowledge of the features or relationships which are influencing the data analyzed. Instead, these are inferred by the ANN through a process of iterative learning. This only requires that the set of inputs and their corresponding outputs is known ex ante, but not their actual relationship. During learning, ANN models process a data set designated for training and utilize an algorithm to adjust their connection weights so that their outputs converge closer to the desired values. While there are many strategies documented in the literature, the most popular learning algorithm is backpropagation (Rumelhart, Hinton, and Williams 1985). Backpropagation is a strategy which adjusts network connection weights using the derivative of the transfer function. The information during learning flows in the opposite direction to the flow observed during processing. This begins at the output layer with a comparison between the ANNs current output and the target output known ex ante. This is used to calculate the deviation between the two also known as error. The derivative of the transfer function is then used to make adjustments to connection weights further down the network, until all the connections are updated. The new information learned is incorporated into the connection weights. The backpropagation algorithm is used for the purposes of the present experiment.



Figure 19: Information flows within Artificial Neural Network

There are many reasons why ANNs may be a suitable technique for carrying out the present experiment. The evaluation of market activity over a short period of time can be seen, in essence, as a pattern recognition exercise. Similarly, the existing commonalities immediately preceding an episode of HFT activity, if present, are not known to the researcher ex-ante. However ANNs do not require such knowledge, as long as all the necessary raw data is fed into the model. Finally, the question of whether a certain window of activity is immediately preceding an HFT episode, can be restated as a Boolean problem, with 1 denoting a period preceding HFT, and -1 any other. The narrower focus of the present experiment is on Quote Stuffing and is specifically looking for a graphical pattern in quote updates immediately prior to the episode of HFT activity. The data generated in the first section of the present research is crucial as input sample for the ANN experiment.

One significant challenge when analyzing two-dimensional data points using ANN models is posed by what is known as the curse of dimensionality. This is a catch

all phrase for many diverse issues arising from the problem of representing two dimensional features in a format suitable for ANN processing. There is a significant body of literature detailing alternative strategies for dealing with this set of issues. For the purposes of the present research, a simplistic approach is adopted, based on 2-D image processing strategies (Egmont-Petersen, de Ridder, and Handels 2002). Each window of quote updates examined, is seen as a 2 dimensional area in time and price. This is further segmented into a number of sections of equal area. The exact granularity of the division along the X and Y axes is determined by the researcher and can be set within the ANN suite of the Trident tool. For the purpose of the present research a granularity of 4 in Price and 5 in Time is selected, yielding 20 segments of equal area. Once these regions are determined, the number of quote update events falling within each segment is estimated, and calculated as a fraction of the total number of quote updates in the time window examined. The end result is an array consisting of 20 fractions denoting relative event density, which sum up to 1. This approach is very similar to the one used in image processing, where images are segmented into areas and pixel counts are performed in each segment to transform the shape of the image into digital form.



Figure 20: Initial segmentation of a chart sample. For demonstration purposes granularity is set to 4 in both dimensions

1	2
3	4

Figure 21: Representation of the 4 input nodes corresponding to the 4 chart

segments	
1 N = 12/30 = 0.4	2 N = 8/30 = 0.267
3 N = 4/30 = 0.133	4 N= 6/30 = 0.2
52.16	09-00-59-523408

Figure 22: Event Density within each region. There are a total of 30 events (quote updates) over the sample period

Figure 23: Representation of the result by filling each region with a % pf black color in accordance with the event density calculated

The resulting set of inputs is readily processed by an ANN model. For the purposes of the present research a training sample of 638 observations is used, with 319 windows immediately preceding a previously detected Quote Stuffing episode, which are assigned a target value of 1, and 319 randomly selected alternative samples which are assigned a desired output value of -1. The ANN models are used to process 600 iteration of the training dataset, and once this is accomplished, a final holdout sample consisting of 50 periods with a target value of +1 and 50 periods with a target value of -1, is used for evaluation purposes.

6 Results and discussion

6.1 HFT strategies and their impact

At a first glance, the present analysis appears to be successful in achieving its main objective of HFT detection. Throughout the data samples examined, a total of 484 potential HFT events are identified, among which 372 episode of Quote Stuffing and 112 episodes of Momentum Ignition. While there is no possibility of verifying these findings within the scope of the present research, some of the patterns observed seem to closely match previous findings in the literature. This may indicate that the detection techniques utilized are appropriate. A further breakdown by time frame show the window sizes used to detect each observation. For Quote Stuffing the majority of events occur within the bigger time scale examined (10s), whereas Momentum Events appear to be more transient, with the majority occurring within 30s, followed by 60s and 90s, as shown in figure 24.

Furthermore, a breakdown within the group of Quote Stuffing events shows that the distribution by strategy subcategory between In-Ask, In-Bid, and Combined are similar in terms of occurrence. A conjecture suggests that Quote Stuffing is used to manipulate the market environment in order to provoke institutional investors to trade. Therefore, trading activity during episodes of Quote Stuffing is expected to be adversely affected by the HFT strategy. For instance, (Tse, Lin, and Vincent 2012) note that during quote stuffing events, the trade prices tend to move in the direction of the stuffing activity i.e. increase when stuffing occurs in-ask or decline when it is happening on the bid side. These effects should be observable in the data, and are therefore tested for. Further results indicate that a majority of ask and bid events are accompanied by trading activity. Furthermore, the characteristic pattern expected is confirmed by using the Implementation Shortfall metric. When IS drift lower to negative values during the time window examined, this is an indication of declining prices. Similarly, as it increases and remains positive, this is an indication of prices rising. Another event is shown below:

It appears that the majority of Quote Stuffing episodes accompanied by trading exhibit the characteristic pattern. These results are especially important since they confirm previous observations. Also they confirm the potential indirect impact of the strategy on market prices. In the process of analyzing the results, a new characteristic pattern is observed which seems to link quote price volatility with Level 1 quoted depth volatility. A frequent observation during Quote Stuffing is that quote updates which narrow the quoted spread appear to be associated by a significant decrease in Level 1 quoted depth. This pattern is very pronounced and may have important implications for correctly interpreting the market impact of HFT activity. Examples are presented below:



Figure 25: Apple quote stuffing in-bid and The level 1 depth at the same interval





Figure 26: Bund stuffing in-bid / combined and Level 1 depth

While these results could potentially be caused by trading activity depleting quoted depth in the order book, the characteristic pattern is also similarly observed during episodes which involve no trades at all. This suggests that the change in volume levels is due to new quotes being posted rather than old ones being depleted. Additionally, these newly posted orders are characterized by very low quantities offered, which is another evidence of HFT activity. One of the most significant impacts of the increasing presence of HFT on markets is a steady decline in the average trade size. This result may be interpreted as indirect proof that the events observed may be caused by the involvement of low latency traders. The observation of small orders being posted and cancelled rapidly over a very short period of time fits the patterns expected of HFT market makers. Moreover, this also lends support to the argument that liquidity provision by HFT may be transient in nature. Finally, this pattern is also consistent with the technique of pinging, since the orders posted narrow the spread and may be intended to entice block traders into execution.

A final pattern is observed at the event level which confirms the intuition that the quoted spread also experiences volatility particularly during one-sided (in-ask or in-bid only) Quote Stuffing episodes. This finding seems to suggest that while HFT seems to provide liquidity through posting small but competitive orders which initially narrow the bid-ask spread, their rapid cancellation increases the volatility and may actually increase trading costs over the long run or even introduce an additional risk-premium for traders. This could potentially offset some or all of the benefits of added liquidity by HFT market makers. A thorough investigation of this hypothesis is beyond the scope of this paper, and may serve as a suggestion for future research in the area of asset pricing (Pastor and Stambaugh 2001)

17:05:35 73



Figure 27: Apple in-ask stuffing and spread volatility



Figure 28: Apple in-bid stuffing and spread volatility

An examination of the intraday patterns of the D-ratio values reveals characteristic peaks around the beginning and the end of the regular trading hours from 9:00 a.m. till 16:00. This was observable across all the three assets as shown below in the figures:



Figure 29: Intraday patterns of D-ratio - Bund



Figure 30: Intraday patterns of D-ratio - APPLE

An examination of the distribution of D-ratio values seems to strongly suggest that these follow a Chi-squared distribution, characterized by a fat tail. This is a finding which warrants further investigations and could potentially lead to a more formalized quantitative method of detecting HFT activities or their probability of occurrence.



Figure 31: Bund combined stuffing D-ratio histogram

Momentum Ignition events are observed on time frames of 30, 60 and 90s. The total size of the price movement for each even is recorded in basis points, as well as the potential volume traded by the HFT traders. The method used for this is an approximation based on the volume observed during the initial volume peak and the secondary volume peak. It is assumed that the strategy practitioners use the initial increase in volume to position themselves accordingly, and once they succeed in instigating a price move, unwinding their position during the secondary volume peak. This is a strong assumption, but the aim of the analysis is simply to provide a potential estimate of the direct economic result derived by the practitioners of the strategy. Future research could refine the methodology used in this estimation, or obtain confirmation of the results from industry.



Figure 32: Momentum Ignition Patterns - Apple

While the average return per event observed in the sample is 23,09 bp, this number is significant when considered within the context of the extremely short time frames of its occurrence between 0,5 and 1,5 minutes. The largest observed relative return in the sample is 106,7 bp, during an event on the Bund futures market. Using the methodology outlined above, the potential gross profit generated during this is EUR 82692,5. This is calculated via the below formula:

$$GrossProfit = Quantity traded * Indicative Price * \frac{Returninbp}{10000} * TickSize (15)$$

Indicative price is a rough indication of the relevant assets price. For Bund futures, this is assumed at a constant EUR 155, for Apple shares at \$113, and for the USO ETF \$17,36. The tick size of the Bund Futures contract is for a nominal value of EUR 10 per 0.01 change in price.

An examination of the intraday D-ratio values chart for Momentum Ignition reveals a similar pattern to the one observed for Quote Stuffing. The pattern observed for Bund contracts is different than the one for the exchange traded assets in the US. This is due to the longer regular hours trading session on EUREX, as opposed to the market hours observed on US exchanges, where both Apple and USO are listed. A start of trading, end of trading and mid-day peaks are observable for most assets.



Figure 34: Intraday patterns of D-ratio - USO

The distribution of the D-ratio values appears to follow the Chi-square distribution pattern observed for Quote Stuffing, but with even longer fat tails and greater skewness, especially for Apple. This suggests granularity in the data, and the presence of extreme outliers as could expected during the episodes of HFT activity. This also provides additional basis for a further future quantitative research on this distribution, which could shed more light on the probability distribution of HFT activity occurrences. The total profits generated by the Momentum Ignition strategy over the sample studied amounts to about \$ 5,25million. A glance at the breakdown of event count by day of the week, reveals no particular pattern, although it seems to suggest that midweek days may tend to contain higher HFT activity.



Figure 35: USO D-ratio histogram



Figure 36: Bund D-ratio histogram

6.2 Artificial Neural Networks Experiment

The ANN experiment is carried out as per the methodology detailed above. This is a simplistic approach, aimed mainly at providing a proof of concept and guidance for further future research. The results seem to suggest that some of the models examined may have forecasting potential. Nevertheless, there are significant caveats with the approach used in this project. A typical problem when using ANN tools, is the large number of potential specifications, and the sensitivity of results to the network architecture used. Small changes to topology may yield extreme differences in performance. Therefore prior selection of the optimal topology is extremely difficult, especially when there is little or no knowledge of the relationship being studied ex-ante. Furthermore, ANN models are very sensitive to the quality of input data. For instance, the present approach focuses exclusively on graphical patterns observed on charts of market activity. The model isnt supplied data about the actual price levels or time frames examined, neither of global long run statistics, such as typical average volatility and volume levels, or other indicators. Similarly, the 20 regions defined in the current experiment may be an insufficient resolution to fully capture the relevant graphical patterns. The experiment detailed here could easily be expanded to include additional information. However, care should be taken to limit the number of input variables, as too many could lead to introducing noise in the data, or to over-fitting. A similar tradeoff is faced when adjusting network parameters such as the learning rate and the learning momentum, as well as the number of neurons in the hidden layers. Adjusting all of these parameters may potentially lead to faster convergence, or the inference of additional features from

the data, however as a downside these may lead to focus on spurious patterns, overfitting or even complete failure to achieve convergence. The table below summarizes the results of several alternative specifications, along with the basic parameters and topology used. Forecasts of negative output, are taken as signaling a result of -1, while positive outputs are seen as signaling +1. While it is difficult to benchmark precisely the forecasting performance observed, a simple rule of thumb is to check whether the forecasts add any incremental value to a nave forecast of 50%.

7 Conclusion

This paper attempts to provide a fresh empirical perspective on the topical area of High Frequency Trading and has four main objectives. A novel empirical methodology is introduced aimed at detecting the presence of two commonly encountered HFT strategies in financial market Quote Stuffing and Momentum Ignition. Further analytical tools are used to assess their performance, impact on the wider market and empirical implications for future research. An experiment using ANNs is carried out to test potential forecasting techniques which may be used to detect these phenomena. All of these instruments are implemented via a proprietary analytical tool developed in the C++ language which has the merits of high performance in processing large data sets, and a lightweight user interface. Section 1 introduced briefly current topics relating to HFT and their general implications for finance academia and industry. Section 2 introduces in detail different HFT strategies, their objectives, implementation mechanisms and potential impact. Section 3 introduces the data and analytical instruments used in the current project. Section 4 presents the results of this analysis and discusses the most important findings and their implications. There are several important weakness of the present project. The main one is the lack of ability to validate the findings presented. Without access to additional data from the industry, it may not be possible to conclude with certainty whether particular market activities can be attributed to HFT or not. Furthermore it is possible that some of the analytical tools used are inappropriate or not calibrated correctly to deal with the objectives of the project. Finally during the course of the current project, several important findings emerged which may serve as a stepping stone for future research. Some these are the spread and volume patterns observed during Quote Stuffing activities, the direct economic performance of Momentum Ignition and the nature of the underlying distribution followed by the D-ratio indicator. The ANN experiment, while serving as a proof of concept, could be vastly improved by introducing additional variables, improving network topology and calibrating its parameters accordingly. Nevertheless, the current project provides some useful analytical tools which can hopefully be of help for other researchers in the field of High Frequency Trading.

References

- Biais, Bruno, and Thierry Foucault. 2014. "HFT and market quality." Bankers, Markets & Investors 128:5–19.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas. 2015. "Equilibrium fast trading." *Journal of Financial Economics* 116 (2): 292–313.
- Bishop, Christopher M. 1995. *Neural networks for pattern recognition*. Oxford university press.
- Egmont-Petersen, Michael, Dick de Ridder, and Heinz Handels. 2002. "Image processing with neural networks review." *Pattern recognition* 35 (10): 2279–2301.
- Foucault, Thierry, Johan Hombert, and Ioanid Roşu. 2015. "News trading and speed." *The Journal of Finance*.
- Kyle, Albert S. 1985. "Continuous auctions and insider trading." *Econometrica:* Journal of the Econometric Society, pp. 1315–1335.
- Lewis, Michael. 2014. Flash boys: a Wall Street revolt. WW Norton & Company.
- OHara, Maureen. 2015. "High frequency market microstructure." Journal of Financial Economics 116 (2): 257–270.
- Pastor, Lubos, and Robert F Stambaugh. 2001. "Liquidity risk and expected stock returns." Technical Report, National Bureau of Economic Research.
- Patterson, Scott. 2013. Dark pools: The rise of the machine traders and the rigging of the US stock market. Crown Business.
- Rijper, Thierry, Willem Sprenkeler, and Stefan Kip. 2010. "High frequency trading." Position Paper. Optiver Holding BV, Strawinskylaan 3095:1077.
- Rumelhart, David E, Geoffrey E Hinton, and Ronald J Williams. 1985. "Learning internal representations by error propagation." Technical Report, DTIC Document.
- Tse, Jonathan, Xiang Lin, and Drew Vincent. 2012. "High frequency tradingmeasurement, detection and response." Credit Suisse, Zürich, Switzerland, Tech. Rep.
- Turbeville, Wallace C. 2013. "Cracks in the Pipeline, Part Two: High Frequency Trading." *Demos, March.*