July 2011

High Frequency Trading

An Overview

José S. Penalva[1]
Universidad Carlos III, Madrid

SUMMARY This document provides a definition of High Frequency Trading, an overview of its position in equity markets, and a discussion of the possible implications for the provision of liquidity, financial market stability, financial innovation, and price stability.

---

[1] This document has grown out of many discussions and debates with Alvaro Cartea whose opinion has naturally greatly influenced my own but all mistakes and opinions expressed herein are mine. I also like to thank Albert J. Menkveld for carefully reading an earlier version and for his useful and thoughtful comments. This paper was written while visiting the Bank of Spain—I would like to thank them for their hospitality. The views expressed in this paper are those of the author and do not necessarily reflect those of the Bank of Spain (Banco de España).

INDEX

1 Introduction

This document provides a definition of High Frequency Trading [HFT], an overview of its position in equity markets, and a discussion of the possible implications for the provision of liquidity, financial market stability, financial innovation, and price stability.

The document is organized in sections: Section 2 provides the definition of HFT strategies, as well as a description of its relationship with other trading strategies. Section 3 considers what are likely sources of profitability for HFT strategies, while section 4 overviews recent equity market developments that are associated with the growth and increased use of HFT strategies. Section 5 looks at the possible impact HFT strategies could be having on the everyday activities of equity markets, and section 6 discusses major issues in equity markets and the possible role of HFTs: liquidity, market stability, financial innovation, and price stability.

2 Trading, Algorithms, and High Frequency Trading

High-frequency traders (HFTs) are firms, trading desks, and individual traders whose primary activity revolves around the execution of high-frequency trading strategies. As the discussion is centered on the use and impact of this particular type of trading strategies, it is useful have a unified framework with which to address issues relating to trading strategies in general. This framework defines a trading strategy as a plan to access financial markets that can be decomposed into five stages:

- Design and planning
- Initial execution
- Holding period
- Final execution
- Evaluation

The first stage consists of identifying the objective of the strategy, as well as identifying and planning all the steps that need to be taken in order to achieve the desired objective. The strategy includes contingency plans and estimates of profits and losses, calculated to ensure the strategy is executed effectively. The second stage, starts at the time when the strategy is put into action and a portfolio of assets is constituted. This portfolio is held over a certain period of time, during which the portfolio is evaluated and adjusted, until the time arrives for final execution, when the portfolio is liquidated. In the final stage, the strategy's effectiveness in achieving the desired objective is evaluated.

For the purposes of this document it is convenient to associate a trading strategy with the completion of a single process starting from the first stage of design and planning, and ending with its final evaluation, although for the most part we will focus on the intermediate three stages: from initial to final execution (including the signals that are used to trigger execution). Needless to say, in practice the same design is applied countless times, and evaluation and execution of trading strategies is done continuous and simultaneously for multitudes of strategies in different stages–to the point that market execution for some strategies is avoided when different strategies generate matching trades (internal crossings).

High frequency (HF) trading strategies are characterized by their short holding periods, mostly measured in seconds, and by their frequent turnover and small asset positions. High Frequency Traders (HFTs) aim at holding no assets overnight and having a large capital turnover, which requires the completion of a very large number of trading strategies every day, every hour, even every minute. Other market participants have different strategies (we follow the characterization in Kirilenko et al. [2010]: normal "slow" market makers have (much) lower turnover which means they trade less often. They also dislike holding inventory but are more willing to hold it overnight. Fundamental traders are characterized by significant trading volume and the accumulation of significant asset positions by the end of the trading day. The remaining traders, usually those that trade very few assets are regarded as noise traders.

In terms of speed of turnover, Kirilenko et al. [2010] estimate, that having bought a block of shares, HFTs will sell half of them within two minutes (on average). But, nothing captures better the nature of HF trading as the following quote from the SEC report [2010]: "16 (out of over 15,000) trading accounts that were classified as HFTs traded over 1,455,000 contracts [of E-mini S&P500 futures] on May 6, which comprised almost a third of the total daily trading volume. Yet, net holdings of HFTs fluctuated around zero so rapidly that they rarely held more than 3,000 contracts long or short on that day." [p. 15] To put some perspective on the number of contracts traded, consider that the average size of a single trade for fundamental traders was between 9 and 25 contracts and almost 7,000 noise traders traded each less than 9 contracts in total that day.

Some people use the terms "HF trading" and "algorithmic trading" interchangeably. This is wrong, as algorithmic trading is defined as "the use of computer algorithms to automatically make certain trading decisions, submit orders, and manage those orders after submission" Hendershott et al. [2010]. In particular algorithms are used in various ways by market participants, not just HFTs, in all stages of their trading strategies: to analyze data, generate trading signals, optimize trade execution, monitor and evaluate portfolios, etc.

Nevertheless, in order for an HF strategy to be successfully implemented very quickly and very often, it is essential that it completes its three main central stages (from initial to final execution) very quickly. Thus, HF strategies are intimately associated with algorithmic (algo) trading at every stage. HF strategies are typically designed via a data and computing intensive analysis of historical trading. They are planned and implemented through an automated computer system. They are executed automatically on the basis of computer generated trading signals. Their execution, maintenance, evaluation, and adjustment is done automatically on computers. Even their final execution is automatic and triggered by computer generated signals, and all results are processed and presented through an automated system.

HF trading strategies are implemented by a trader's "trading system". Once implemented, human intervention in HF trading strategies is limited to determining "whether the system is performing within prespecified bounds, and if it is not, whether it is the right time to pull the plug" [Aldridge [2009] p. 5]. p5].

Thus, HFTs and algo trading are intimately linked and it is very hard to distinguish which part of observed algo trading is HFT and which is not. This problem is particularly important when interpreting the results of empirical studies such as those of Brogaard [2010] and Hendershott et al. [2010], and leads to a lot of confusion in less specialized publications and the public debate.

3 HFT Profits

Prerequisites for the existence of HFTs are fully electronic (and fast) trading platforms and low transaction costs. The profits (and losses) from any single trading strategy implemented by a HF trader are intrinsically limited by the short holding period and the market circuit breakers. Historically, prices have moved relatively little for the short duration of HF trades. And exchanges have automatic "safety" measures that impose limits on short-term price movements. With such built in limits, a small tax or exchange fee can wipe out the profitability of HF trading as we know it. Nevertheless, these limits can be misleading given the speed and volume of HF trading. The accumulated effects of HF trading strategies can lead to very large profits (and losses), and influence market wide events (as during the Flash crash).

HF traders can be classified into three groups: sell-side institutions (e.g. Goldman Sachs), hedge funds (e.g., Citadel or Renaissance) and pure high-frequency-trading firms (e.g., Getco). They are (understandably) very secretive, so that only they know where their profits come from. Despite this, one can identify broad categories of strategies that are suitable for HF trading:

- Exploiting the spread
- Transparency (limit order book) arbitrage
- Cross-market arbitrage
- Event arbitrage
- Statistical arbitrage

Market making behavior has been extensively analyzed in the financial literature, and a fundamental source of profits comes from managing asset inventories while making money from the spread (the positive difference between the best sell offer and the best offer to buy). Holding inventories of assets is costly in financial terms. If a market maker buys an asset and immediately resells it, he makes money (the spread). But having bought an asset, a trader may not be able to sell it immediately which puts him at risk. In particular, he risks that by the time he finds someone to sell the asset to, the price of the asset has changed. This price change may be temporary, either up or down but neutral on average, so he requires some compensation for this risk. But, the market maker also faces a second type of risk, namely, that he is trading with informed investors who know where the market is going. Thus, he may be accumulating assets whose price is falling (and making a loss on the transaction). The spread is supposed to compensate market makers for accumulating these assets and absorb these risks (that is, for providing liquidity) even when the price is falling, something they are obliged to do—official market makers are required by the exchanges to always provide liquidity. On the other hand, HFTs can choose when to provide liquidity and when not to. HFTs, because of their speed, have better control of when their posted offers are going to be executed, and are thus able to extract some of the spreads that would normally go to regular market makers.

Transparency arbitrage exploits the fact that by placing a limit order, you give the market a free option. If you place a limit order to sell at a price *P*, you are saying to the market "if the value of the asset is at or higher than *P*, I will sell you the asset at a price of *P*". As Shiller [2010] puts it in his class notes: "Placing a limit order is being a sitting duck. If more in-the-know investors realize that the stock price is higher than in your sell order, they will buy from you at your price. In an expected value sense, your orders can be filled only in this unfortunate case." Exchanges have moved towards increased market transparency, and in particular, towards increasing the amount of information they provide about outstanding orders pending execution. These outstanding orders are by definition limit orders, and one of the aims of HF strategies is to filter through that information and exploit those free options.

Arbitraging across markets is a long-practiced strategy that consists of making money from differences between prices for the same product at two different trading locations. When a broker-dealer receives an order to execute a trade from a client, he has many options at his disposal: The order can be crossed internally, as the broker matches this order with another received from a different client; the order can be entered manually into one of the exchanges (for example, NYSE or NASDAQ); or, the broker can launch a trading algorithm that will explore different trading venues looking for the best execution (these trading venues (ATS, ECN, "dark pools", etc are reviewed in more detail below). The existence of multiple trading venues is a fertile ground for fast and efficient algorithms to find minor price differences to exploit.

Many HFTs are continuously monitoring news announcements as a source of profits. HFTs do not have the time to determine what an asset's "objective" value will be in response to a particular piece of news and build positions around news events. On the other hand, HFTs understand very well the price patterns surrounding certain types of news and can exploit systematic price movements around these announcements, such as pre-trade information leakage (prices anticipate the direction of the stock movement in response to news) and post-trade overshooting (prices tend to exacerbate the magnitude of the stock movement just after the news announcement).

Finally, there is a great deal of literature written on statistical arbitrage. Essentially statistical arbitrage amounts to the usual notion of arbitrage–making profits from unjustified price differences–only it exploits differences that arise, not between the same object in different markets or contractual form, but from similar objects, where similarity is established in the sense that there are persistent stable statistical relationships between two objects. Borrowing an example from Narang [2009]: suppose you think two stocks should behave similarly (for example Exxon Mobil and Chevron). If their prices start to diverge then one would expect that the difference will be reversed shortly. Within an HF trader's time scale, this amounts to something as simple as spotting an uptick in Exxon Mobil and revising up offer prices on outstanding sells combined with an increase in the aggressiveness of buys for Chevron, or it could involve more sophisticated strategies with prolonged holding periods.

4 Market Developments (US Equities)

From a historical perspective, HF trading is a byproduct of the transition from manual to automated trading in all exchanges, and the associated reduction in direct costs of execution (both financial and in terms of time). Since the 1990s we have seen a number of significant changes: the introduction of Order Handling Rules (adopted in 1996), the change from quoting prices in sixteenths of a dollar to decimalization on the NYSE (January 29, 2001), and the adoption of the controversial Regulation NMS (2005, January 1, 2006) designed to adapt equity trading rules to the new electronic trading environment.

Since 2005 we have seen a substantial change in US Equities markets as reflected by the following numbers (SEC [2010]):

- Percentage of NYSE-listed stocks' volume executed at the NYSE: 79,1% (2005) vs. 25.1% (2009)
- Average speed of execution for small, immediately executable orders (NYSE): 10.1 seconds (2005) vs. 0.7 seconds (2009)
- Consolidated average daily volume in NYSE: 2.1 billion shares (2005) vs. 5.9 billion (2009)
- Consolidated average daily volume of NYSE-listed stocks: 2.9 billion shares (2005) vs. 22.1 billion (2009)
- Consolidated average trade size in NYSE-listed stocks: 724 shares (2005) vs. 268 shares (2009)

Furthermore, stocks (as well as other assets) now trade in a large number of very diverse trading

Registered exchanges:			
	NASDAQ	19.4%	
	NYSE	14.7%	
	NYSE Arca	13.2%	
	BATS	9.5%	
	NASDAQ OMX BX	3.3%	
	Other	3.7%	
Total Exchange			**63.8%**
ECNs:	2 Direct Edge	9.8%	
	3 Others	1.0%	
Total ECN			**10.8%**
Total Displayed Trading Center			**74.6%**
Dark Pools:			
	Approximately 32	7.9%	
Broker-Dealer Internalization:			
	More than 200	17.5%	
Total Undisplayed Trading Center			**25.4%**

venues. In particular, the SEC has published a concept release on equity market structure where we find the following table describing the distribution of trading volume across trading centers:

Before we proceed let us review in a bit more detail what are these different trading venues. A broker can try to execute an order internally (against its own capital or a balancing order from another client), it can execute the order electronically, or he can personally negotiate a deal. The primary method is electronic execution, and this can occur at a regulated exchange (such as NASDAQ, NYSE, a regional exchange, or, since July 2010, DirectEdge). It can also take place over an ECN (Electronic Crossing Network) or at a "dark pool". Most trading venues that are not regulated exchanges, including dark pools, are ATS (Alternative Trading Systems). An ATS differs from a regulated exchange in that an ATS "is not required to file proposed rule changes with the Commission or otherwise publicly disclose its trading services and fees. ATSs also do not have any self-regulatory responsibilities, such as market surveillance" (SEC concept release on Equity Market Structure). In Europe, the primary alternative to a regulated exchange is a MTF (multilateral trading facility) and is regulated according to the MiFID (Markets in Financial Instruments Directive).

The primary difference between ECNs and dark pools is that ECNs offer trading services analogous to those offered by exchanges (so analogous that Direct Edge has filed for and obtained full exchange status). ECNs (and exchanges) provide pre-trade information (such as best-price quotations) as well as post-trade reporting of trades as they are executed. Dark pools, on the other hand, do not provide pre-trade information and provide the (minimal) post-trade information required by law (RegNMS). Some of the main dark pools for US equities are Crossfinder (Credit-Suisse), SigmaX (Goldman Sachs), Liquidnet, POSIT, and BIDS. Liquidnet and POSIT are independently owned, while BIDS is a joint venture between major investment banks and NYSE. The main dark pools in Europe are Chi-X, Turquoise, BATS, Instinet, Liquidnet and Smartpool. On April 04, 2011, the UK clearing house EuroCCP and Turquoise have announced a new service for Spanish stocks. Nevertheless, hidden trades are posted in many exchanges, not just dark pools. Regular exchanges (such as NASDAQ) give traders the option to post hidden (or partially hidden) offers, that are not displayed prior to trade execution.

Returning to trading volume in US equity markets, the table above provides an overview of how fragmented the market for US equities is. We can see that a quarter of traded volume in equities takes place outside regular exchanges. The other three quarters of traded volume takes place within highly automated trading systems that can offer extremely high-speed, or "low-latency," as well as order responses and execution which can be as low as less than 1 millisecond. Furthermore, many exchanges offer individual data feeds that deliver information concerning their orders and trades directly to customers, as well as the possibility of co-location. Co-location is the right for a customer of the exchange to place their servers in close proximity to the exchange's matching engine (which helps reduce latency). Also, for some time broker-dealers allowed some of their clients to use "naked access" (unfiltered and unchecked access to the exchange). These customers could connect directly to submit their orders without any checks or supervision. This practice has been recently forbidden by the SEC and the ban is currently being implemented.

Exchanges currently offer their customers many options as to how they can execute trades that go well beyond the traditional dichotomy of limit and market orders. Today, some of the things traders can do are: specify which liquidity pools to check, whether to display the whole order on the book or only a fraction (reserve orders) or nothing at all (non-display orders), whether to "peg" their orders to the market price (post offers with prices relative to the NBBO, ask, bid or mid-point), or, they can post discretionary orders (orders that are displayed on the book at one price while passively trading at a more aggressive discretionary price).

Within this context we have lived through some interesting financial episodes, the most remarkable ones being the rise and fall of credit derivatives (credit default swaps and collateralized debt obligations), and the Flash crash of 2010. The latter is strongly associated with high-frequency trading and is very well documented. In fact, the analysis of the audit trail during that episode included in two documents, the SEC [2010] report and Kirilenko et al. [2010] provides the best and probably the only close detailed look at broad high frequency trading that is publicly available, as they had access to the identity of the participants in every trade.

The events of May 6, 2010, in equity markets do not represent an entirely isolated incident, as demonstrated by the cocoa flash crash on March 1, 2011: between 10:00 and 10:30 a.m. the price fell by 12.5% in less than a minute, and then it recovered (starting around 3,700$ and ending up around 3,610$). Also, the price of sugar saw a sudden (and temporary) price crash as the price dropped 6% on February 3, 2011, in one second. These events, like the Flash crash, are associated to HF trading but no detailed analysis exists like that for the Flash crash that would allow us to reach any certain conclusions.

5 The Impact of HFTs

An outside observer has to deduce what is going in equity markets from limited information. The first problem for an outside observer is that HFT behavior is hard to disentangle from that of other algo trading strategies. Making this distinction is very important because any policies aimed at HFTs will need to cut through the Gordian knot that ties them together with other algo trading strategies. Nevertheless, a number of scholars have identified patterns in the data that are associated with HFT-like behavior.

HFTs rely on speed. Because of their speed, HFTs are the first to act and react. Thus, as we saw above, they are there to sell to those who want to buy and to buy from those who want to sell, acting as traditional market makers and providing "liquidity" to the market. We see the effect of this speed in the reduction in execution time which is now less than one second.

HFTs are fast, very fast. This implies that they can change their trading positions in the order book (post, change and cancel limit orders, send market orders, etc) very quickly, and their movements appear in prices and the order book. The traditional idea that limit orders sit there "patiently supplying liquidity" is challenged by the speed with which many limit orders are being cancelled. Even in October 2004 Hasbrouck and Saar [2007] report that more than one third of limit orders were cancelled within two seconds on INET (NASDAQ). Hendershott et al. [2010] in their study of the effect of the introduction of automatic quote dissemination in the NYSE conclude that "algorithms can easily take into account common factor price information and adjust trading and quoting accordingly".

And speed is important not just to act, but to process information. This gives HFTs an informational advantage. Hendershott and Riordan [2009] find that algorithmic trades have a larger permanent price impact than those of their human counterparts and exhibit behavior consistent with "being able to better disguise their trading intentions." Hendershott et al. [2010] point to the presence of "algorithms ... designed to sniff out other algorithms or otherwise identify order flow and other information patterns in the data". Kirilenko et al. [2010] similarly conclude that "possibly due to their speed advantage or superior ability to predict price changes, HFTs are able to buy right as the prices are about to increase."

We also observe large price spikes in short time spans in a way that was not there before. The Flash crash is unique although we have had very active HFTs during the last five years. This could be due to a temporary withdrawal or delay in the provision of liquidity from "slow" traders that leaves HFTs trading amongst themselves, generating "hot-potato" effects that can provoke sudden price spikes (as is believed to have happened during the Flash crash—SEC [2010]).

Computers never get tired. They monitor many markets constantly and in parallel. They can immediately spot "unjustified" price differences. As trading costs go down, many more of these price differences are profitable, and exploited extremely quickly. Chaboud et al. [2009]'s observation of the exchange rate market for euro/$, $/yen, yen/euro in 2006-07, conclude that "computers [are] taking advantage of short-lived triangular arbitrage opportunities" as opposed to their human counterparts. Also, computers are systematic. They are not prone to mood swings or flights of fancy, and yet they are observed to flock together more than their human counterparts, at least in the FX market; Chaboud et al. [2009] find that computers are more likely to trade with each other than theoretically predicted in an environment where all traders are randomly matched.

Algorithmic trading and HFTs are part of the trading landscape. Furthermore, the number of firms and activity surrounding the design and implementation of computerized trading systems and execution algorithms is growing. We now turn to the issues this raises.

6 Open Questions on the Role of HFTs

The first issue, and an often repeated claim, is that HFTs provide liquidity. Hendershott and Riordan [2009] find that "while AT [Algorithmic Traders] supply liquidity for exactly 50% of trading volume, AT are at the best price (inside quote) more often than humans. This AT-human difference is more pronounced when liquidity is lower, demonstrating that AT supply liquidity more when liquidity is expensive." HFTs are acting as market makers.

And yet, HFTs cannot (by design) act as liquidity providers over a significant time horizon. Their holding periods are limited by construction. Although they provide liquidity at first, they will quickly turn around and undo their positions, which may involve extracting liquidity from the market. As we discussed above, HFTs (successfully) compete with traditional market makers to act as counterparty to those coming to the exchange to buy and sell assets, making money from the price spread when they buy from one person and sell to the next. The fact that they are there to act as counterparties to a trade (provide liquidity) is a valuable service to the market, but the problem is that HFTs are not willing to hold the asset for long, so if no one comes to buy from them quickly enough, they will actively sell it, which implies they will withdraw liquidity. This helps to understand why we see HF trading to be overrepresented in large cap stocks, where there are plenty of other traders willing to provide that extra liquidity if the HF trader needs to actively close the trade. Furthermore, we must consider that what we observe includes the reaction of other market participants. It is not clear that the liquidity provided by HFTs is not just substituting for that which would have been provided by slower traders.

Other HF strategies also have an indirect effect on liquidity. HF cross-market arbitrage strategies impound information into prices quickly. They quickly spot unjustified price differences and arbitrage them away. This helps keep prices efficient and allows information to travel quickly across markets and different assets ... for better or for worse. While liquidity is plentiful, HFTs keep prices in line and expand the universe of possible counterparties, thereby increasing the possibilities of finding a better match for a trade and welfare. But when overall liquidity dries up in one market (as in the Flash crash), its effect reverberates across all markets very fast.

HFTs have further indirect effects on liquidity, or rather, on slow trader's incentives to provide liquidity. Empirically, "slow" traders provide less liquidity, or rather they provide liquidity less aggressively. From the study of the audit trail of the Flash crash Kirilenko et al. [2010] conclude that one of the lessons to take away is that traders, especially "slow" traders implement "their own versions of a trading pause". This self-imposed trading pause suggests that "slow" traders are cautious and will stop trading if they feel the market is doing something unusual. Consider the slow trader's problem: they face HFTs that are faster to register changes, see more activity than them, process information faster and more systematically, have spent a great deal of resources on programming talent, and can make 100 trades in the time it takes a human trader to read a headline, switch to the trading interface and click on the button to execute a trade. Would you be cautious? And caution is infectious ... so much so that it seems like it played a key role in the momentary liquidity crisis that provoked the Flash crash.

Caution is not the only response. Slow traders are speeding up. "Buy-side" traders (such as investment funds and pension funds) are boosting their own algo trading tools, buying "guerrilla" algorithms designed to optimally slice large orders and execute them over the many liquidity pools, carefully so as to minimize price impact and avoid detection by the HFTs' "shark" algos (algorithms design to identify and exploit large pending orders)–the drastic fall in the average number of shares per trade is a clear indication that such or similar strategies are becoming widespread. There are also "sniffer", "sniper", "smoke" and a host of other similarly named algorithms that suggest a war-like situation. And there are commentators that forcefully argue that traders are in a state of war, an arms race to acquire the fastest, most efficient computer hardware, software, and programming talent. Technological innovation is usually good, but when it is designed to fight zero-sum games, one wonders whether there may be too much of a good thing.

Also, amongst all this accelerating trade, an important question is what happens with prices and the orderly functioning of markets. Some commentators talk about "bait-and-switch" strategies: orders posted at attractive prices that disappear by the time market orders arrive, and are replaced by quite unattractive ones; or strategies pegged to your orders that drive up market prices as you try to buy and lower them as you try to sell. Evidence for the widespread use of such strategies is scant. Furthermore, regulation NMS was designed to limit such strategies, setting up the NBBO (National Best Bid and Offer) to serve as a common source of price information for all regulated exchanges. Nevertheless, even the SEC recognizes that the NBBO is not a perfect solution: in SEC [2010, p. 78] they acknowledge the possibility of arbitrage between exchanges and dark pools that use NBBO as the reference for settling matching trades.

Prices are also increasingly volatile. Equity prices were never very stable, even though they are supposed to reflect the value of the enterprise that issued them. These prices are used as reference points for many purposes outside everyday trading. But, with the large increase in the number of trades, it is not surprising that there should be an increase in the number of price changes. And yet, with the speed at which things are moving in equity markets, one may wonder how useful these prices have become. Picking unique but illustrative circumstances, consider the problem of traders who, at the time of the Flash crash (which originates from S&P500 futures market), sent in market orders to buy equity of blue chip companies, such as Procter & Gamble, and found their orders being executed at prices as low as one penny and prices as high as $100,000 (SEC [2010, p. 1]). Of course, under normal circumstances prices do not fluctuate anywhere nearly as much, but the issue remains as to what the price of an asset is, where by the time you see the ticker on CNBC and you boot up your computer, there has been more trading than would have occurred in two hours, fifteen years ago.

7 Conclusion

With the growth in electronic trading, a new breed of traders has entered the market. These traders build trading systems that do their trading for them. These systems are blazingly fast, and are able to obtain large profits by accumulating tiny ones from a very large number of small trades over the course of a day, day in and day out. They scratch cents from trading on the spread, off the trading book, arbitraging across markets and assets, and around trading news. And they have vocal validators and detractors. The former claim that markets are more liquid, more efficient and there is more competition between exchanges which lowers trading costs without problems from market segmentation. The latter claim that HFTs impose a tax on trading by increasing the costs of execution, that they do not add but only substitute for existing liquidity, that they have provoked a useless technological arms-race, that they threaten the stability of the financial system, and that they have turned financial markets into nothing better than a casino.

We hope to have shed some light on these issues. As we have seen it is hard to separate the Dr. Jekyll in algorithmic trading from the shadowy HF trades of Mr. Hyde. Both are an intrinsic part of modern electronic financial markets. The challenge for society and regulators is to ensure financial markets remain open and competitive, without losing sight of what their primary function should be: to facilitate the movement of financial resources to their most productive uses.

8 References

I. Aldridge. High-frequency trading: A practical guide to algorithmic strategies and trading systems. Wiley, 2009. ISBN 0470563761.

Jonathan Brogaard. High frequency trading and its impact on market quality. SSRN Working Paper, 2010.

A. Chaboud, B. Chiquoine, E. Hjalmarsson, and C. Vega. Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. 2009.

U.S. Commodity Futures Trading Commission, the U.S. Securities, and Exchange Commission. Findings regarding the market events of may 6, 2010. Report, SEC, September 2010.

Joel Hasbrouck and Gideon Saar. Technology and liquidity provision: The blurring of traditional definitions. SSRN eLibrary, 2007.

T. Hendershott and R. Riordan. Algorithmic trading and information. In Working Paper presented at First FBF/IDEI-R Conference on Investment Banking and Financial Markets, Toulouse, 2009.

Terrence Hendershott, Charles M. Jones, and Albert J. Menkveld. Does algorithmic trading improve liquidity? Journal of Finance, Forthcoming, 2010.

Andrei A. Kirilenko, Albert (Pete) S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. The Flash Crash: The Impact of High Frequency Trading on an Electronic Market. SSRN eLibrary, 2010.

R.K. Narang. Inside the black box: the simple truth about quantitative trading. John Wiley & Sons Inc, 2009. ISBN 0470432063.

SEC Concept Release on Equity Market Structure, Release No. 34-61358; File No. S7-02-10

Robert Shiller. Financial Markets. Open Yale Courses, 2010.