High Frequency Trading

Bruno Biais, Toulouse School of Economics & Paul Woolley, London School of Economics[1]

March 2011

Preliminary

Comments welcome

[1] This paper benefitted from the support of the Paul Woolley Research Initiative on Capital Market Dysfunctionality at IDEI Toulouse and the Investment Banking and Financial Markets Value Chain FBF chair at IDEI Toulouse. Many thanks for helpful comments to seminar participants at the SUERF/UPF Conference in Barcelona on December 10, 2010, the conference "*Market Microstructure: Confronting Many Viewpoints*," in Paris on December 6, 2010, and the FBF workshop on algorithmic trading in paris on October 14, 2010. We also benefitted from very useful discussions with Sébastien Cortez from Société Générale, Charles Albert Lehalle from Cheuvreux-CréditAgricole, Hubert Rotivel and Vivien Lévy Garboua from BNPParibas, Arnaud Oseredczuk and Laurent Grillet Aubert from AMF, Jean Tricou from FBF, and Jeremy Grantham from GMO, as well as Christopher Clack, Fany Declerck, Thierry Foucault, Xavier Freixas, Terry Hendershott, Johan Hombert, Pete Kyle, Philipp Hartmann, Augustin Landier, Bert Menkveld, Sophie Moinas, Sébastien Pouget, Xavier Vives, and Pierre Olivier Weill. Remaining errors are of course ours.

High Frequency Trading

Two of the main functions of financial markets are to convey price signals to the economy and to enable agents to reap gains from trade. Financial innovation is useful to the extent that it enhances the efficiency with which markets serve these functions.

High-frequency-trading (hereafter HFT) is one of the major recent innovations in financial markets. It was estimated in 2010 by consultancy *Tabb Group* to make up 56% per cent of equity trades in the US and 38% in Europe.[2] HFTemploys sophisticated computer programs to analyze market data in the search for trading opportunities that may open up for anything from a few seconds to a few hours.[3] Computers then map this information into trading strategies and route orders to market venues, all without direct human intervention.[4]The speed of this process is baffling. The "latency" between the arrival of information at the computer and the execution of the order is of the order of milliseconds, far faster than humans can even register the initial information. High frequency traders compete for speed both by having the most powerful computers, connections and programs and by paying premiums to the exchanges for the privilege of locating their computers as close as possible to the trading venue.

HFT has been the subject of intense public debate and controversy. Some commentators argue that it increases trading volume and liquidity, lowers trading costs and helps price discovery, and is therefore a socially beneficial financial innovation. Others claim it may increase volatility and systemic risk and creates a non-level playing field. Some regulators have expressed concern. For example, SEC Chairman Mary Schapiro said in a speech on September 22, 2010, "...high frequency trading firms have a tremendous capacity to affect the stability and integrity of the equity markets. Currently, however, high frequency trading firms are subject to very little in the way of obligations either to protect that stability by promoting reasonable price continuity in tough times, or to refrain from exacerbating price volatility." But other regulators have shown tacit acceptance or support.

---

[2]Jeremy Grant (Sept. 02, 2010)."High-frequency trading: Up against a bandsaw". *Financial Times*.http://www.ft.com/cms/s/0/b2373a36-b6c2-11df-b3dd-00144feabdc0.html.

[3] Fundamental information hits markets at relatively low frequency. Hence the data analyzed by HFT firms is in its vast majority prices and orders, except when there are news announcements, in which case the competition between computers and humans involves the speed at which they react to the announcement and also the accuracy of their interpretation and analysis of the news.

[4]Chlistalla (2011) offers a description of several typical HFT strategies involving liquidity provision or detection, and statistical arbitrage.

The goal of the present paper is to shed some light on these issues, with a view at informing the policy debate.[5] In doing so, we rely on evidence gathered from discussions with market participants and from empirical studies (Hendershott, Jones and Menkveld (2010), Hendershoot and Riordan (2009), Chaboud et al (2009), Brogaard (2010) and Kirilenko et al (2010)). To interpret the evidence and delineate economic and policy implications, we rely on theoretical analyses of trading – and particularly high-frequency trading – in financial markets as well as on first principles economic reasoning.

The next section offers descriptive information on HFT. Section II presents the motivation for HFT. Section III discusses its costs. Section IV offers an overall assessment. Section V offers some perspectives about likely future evolutions and discusses policy implications.

I) Description

There is a long-term trend towards more and more electronic financial markets. In the 1970s, the NYSE introduced the "DOT" (designated order turnaround) system which routed buy and sell orders to the proper trading post for manual execution. During the 1980's, stock exchanges moved progressively to electronic limit order books, allowing investors to directly route limit and market orders to trading platforms automatically matching them. In the 1990's,electronic communication networks offerednew trading venues that broke the monopoly status of the established central stock exchanges.

As exchanges and trading platform moved towards a more and more electronic model, so did investors. Investing on the basis of algorithms has been widely practiced for, as long as computers were good enough to handle the data.Since the 1980s, quantitative strategies, based on value criteria or price momentum, have been particularly popular with investment funds and prop traders. At the same time, program trading developed to exploit and corrects mispricing between related assets.[6]The last decade saw the development of higher frequency algorithmic-trading. Buy-side investors now extensively rely on algorithms to optimize the execution of their holding strategies. In these trading algorithms, computers determine, as a function of preset parameters and current market variables, the timing, price, quantity and routing of orders. This leads, in particular, to the splitting of orders through time and across market venues, and the strategic choice between market and limit orders. While implemented on a rather short-term basis, these algorithms are designed in view of the relative long-term holding policies of buy-side investors. In contrast, sell-side institutions (e.g. Goldman Sachs), hedge funds (e.g., Citadel or Renaissance) and so called pure-play high-frequency-trading firms (e.g., Getco), have developed trading algorithms involving very short-term holding periods.[7] Such high-frequency-trading (hereafter HFT) is the focus of the present paper.

[5] The Dodd Frank Wall Street Reform and Consumer Protection Act called for an in depth study of HFT (Section 967 (2)(D)).

[6] Harris and Shapiro (1994) and Hasbrouck (1996) offer high-frequency data analyses of program trades and liquidity on the NYSE.

[7] One of the quintessential qualities of HFT algorithms is to be extremely fast. To achieve this, HFT firms rely, in addition to high speed connexions and co-location, on parsimonious, short codes. Also, since the horizon of these algorithms is very short, to achieve parsimony they may dispense from on-line learning features. Yet, at the same as the algorithm is running, the HFT firm may well be simulating the use of competing algorithms, and choose to switch from the current one to another one, if the latter appears to perform better given market conditions.

While the above offers a definition of HFT, for empirical studies it can be difficult to identify which trades and orders stem from HFT. Initial empirical studies relied on proxies for the aggregate level of algorithmic trading (see, e.g., Hendershott, Jones and Menkveld, 2010). More recent studies have been able to rely on databases where orders placed by certain categories high-frequency traders were identified. Hendershott and Riordan (2009) took advantage of a program introduced by the Deutsche Boerse in December 2007, whereby institutions received significant trading fees rebates for trades they could prove stemmed from automated trading. Kirilenko et al (2010), using audit trail data which the CME provides to the CFTC, classify as high-frequency-traders those intermediaries which were among the top 3% in terms of number of orders. Brogaard (2010) uses a dataset distinguishing the orders from 26firms that Nasdaq identified as high frequency traders.[8]Chaboud et al (2009) study data from EBS, one of the two major foreign exchange trading platforms (with Reuters). Traders cane enter instructions manually using an EBS keyboard or via computers directly interfacing with the system. It is the latter type of order which is interpreted as computer generated algorithmic trading in Chaboud et al (2009).

HFT was initially developed in the context of equity markets. Yet, as electronic trading platforms & exchanges developed, it spread to options, futures, ETF's (exchange traded funds), currencies, and commodities For example, on the Intercontinental Exchange, which bought the New York Board of Trade in 2007, open-outcry futures trading has been replaced with computer trading and high-frequency trading is estimated to make up approximately 10% of trading volume (see Gregory Meyer, Financial Times, March 9, 2011, “High-speed commodities traders under scrutiny”).

HFT and more traditional "Quantitative Investing" have a number of features in common, both in the uses to which the algorithms are put and the impact they have on market performance. Both seek to identify and exploit systematic price patterns, such as, e.g., momentum or reversal. HFT, however,concentrates more on the information that price movements might convey about the trading strategies or signals of other market participants. The key difference is the holding period or investing horizon. That of HFT ranges between milliseconds and hours. Their entire positions are closed at the end of each trading day. The type of momentum they focus on is therefore of a much more short-term nature than that exploited by traditional quantitative investing. To illustrate the short horizon of high frequency traders’ holdings, Figure 1 plots their dynamics over several days. Panel A, borrowed from Jovanovic &Menkveld (2010), plots the dynamics on January 30, 2008, of the inventory of the high-frequency trader they study. The figure shows the rapid oscillation of the position of the trader around 0. Panel 2, borrowed from Kirilenko et al (2010) depicts the net position of high frequency traders (as well as transactions prices) in the June 2010 E-mini S&P futures contract over one minute intervals during May 3,4, 5 and 6.[9]

---

[8]These firms have the following characteristics: They engage in proprietary trading, they don’t have customers but trade on their own capital, they use sophisticated trading tools such as high-powered analytics and computer co-location services, they switch between long and short positions several times throughout the day, their orders have a short duration, and they have a low ratio of trades per order.

[9] Another manifestation of short horizon of HFT algorithms is the rapidity with which their limit orders are cancelled. The existence of such rapidly cancelled “fleeting orders” was first documented by Hasbrouck and Saar (2011).

Figure 1, Panel A:

Position of the high frequency trader studied by Jovanovic and Menkveld (2010) , aggregated across Euronext and ChiX, January 30, 2008.

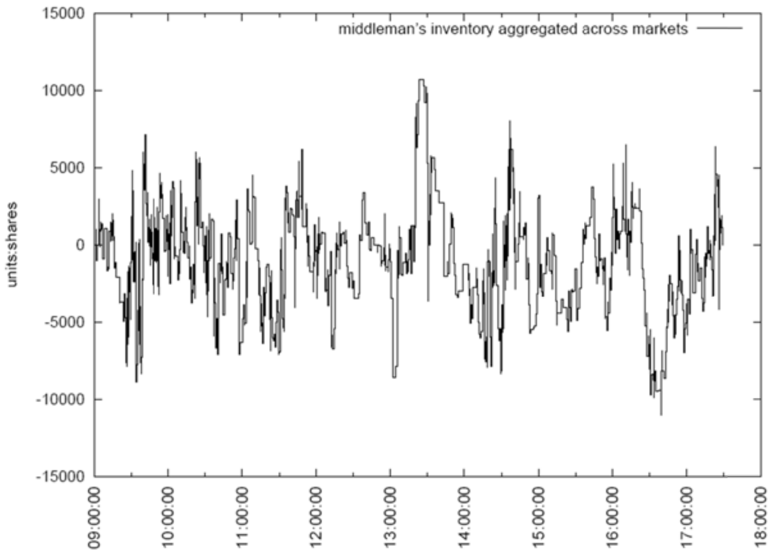

Figure 1, Panel B:

Net position of high frequency traders and transactions prices in the June 2010 E-mini S&P futures contract over one minute intervals during May 3,4, 5 and 6. Kirilenko et al (2010)

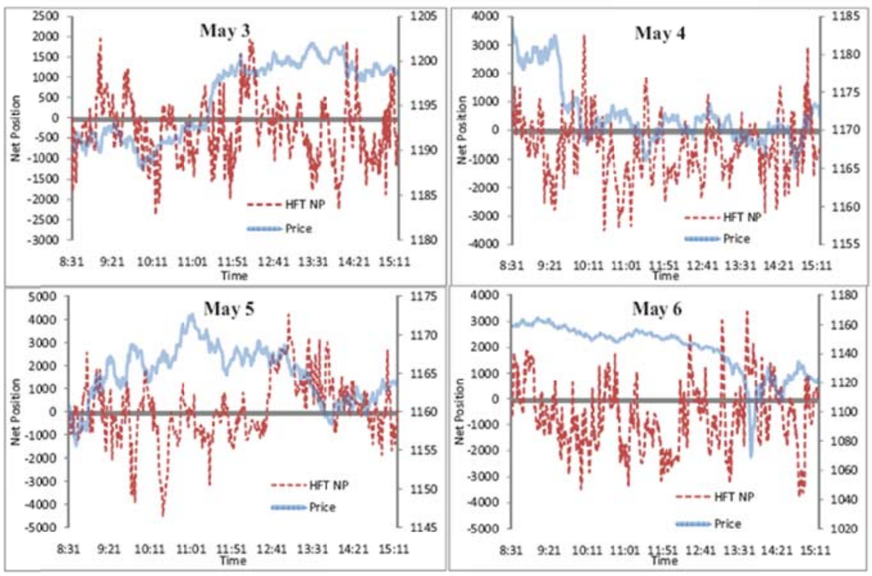

II) Motivation for High Frequency Trading

Trading algorithms collect and process information faster than, and instead of, humans. Why is it useful? The development of algorithmic trading reflects a trend towards automation and computerization which affects all aspects of service industries, and generates productivity gains. Thus, it reduces the operating costs of banks and financial firms. In addition to these cost savings, it can, in some respects, improve the functioning of markets.

II.1) HFT algorithms can help ensure that related assets remain consistently priced

Chaboud, Chiquoine, Hjalmarsson and Vega (2009) offer interesting evidence on this point. They study HFT in the foreign exchange market, focusing on the top two most trade currency pairs (euro-dollar and dollar-yen), as well as the euro-yen cross rate. As the authors note: "In this cross-rate we believe computers have a clear advantage over humans in detecting and reacting more quickly to triangular arbitrage opportunities, where the euro-yen price is briefly out of line with prices in the euro-dollar and dollar-yen markets." Computers are quick at identifying and exploiting these profit opportunities, and by doing so will quickly bring back in line the currency prices.[10]

Algorithmic trading can also help investors realize mutually beneficial trades faster. This is socially valuable to the extent that it enables investors to share risk before it is realized (and thus escape Hirshleifer's effect) and assets to be rapidly transferred from those who value them less towards those who value them more. As discussed blow, trading algorithms can achieve this by i) reducing search costs and delays, and ii) mitigation the traders' cognition limits.

II.2) HFT algorithms can help traders cope with market fragmentation

The trend over the past decade or so has been for the emergence of multiple trading platforms. Figure 2 illustrates for European markets the resulting coexistence between traditional exchanges and new entrants. How did such market fragmentation emerge?

An investor's first consideration will be to find a location that gives the best chance of finding counterparties. There is therefore a centripetal force that causes trading in a security or group of securities to be concentrated in a single location (see Pagano, 1989 a & b). But access to trading in a single location may be regulated or taxed by the owners of the exchange either through membership charges or regulation of the commissions charged by brokers. The advantage of liquidity may therefore be offset by monopolistic pricing, and investors may find it worthwhile sacrificing higher liquidity for lower charges by moving their trading to alternative locations. The formation of new venues thus promotes competition for cheap and effective trading facilities.[11] For this reason, it has been looked on favorably by policy-makers and regulators. The Markets In Financial Instruments Directive (MIFID, in force in Europe since 2007) and the Regulation National Market System (regNMS, in force in the US since 2005), have fostered competition between trading mechanisms, including exchanges and other platforms.

[10] Such HFT is a modern form of standard program trading.
[11] Island was one of the most successful of these new trading platforms. Biais, Bisière and Spatt (2010) analyze how liquidity supply on Island competed with its Nasdaq counterpart.

Figure 2:
Market share (in Euros) of the major equity trading platforms in Europe in 2010 (author's computations, based on FESE data)

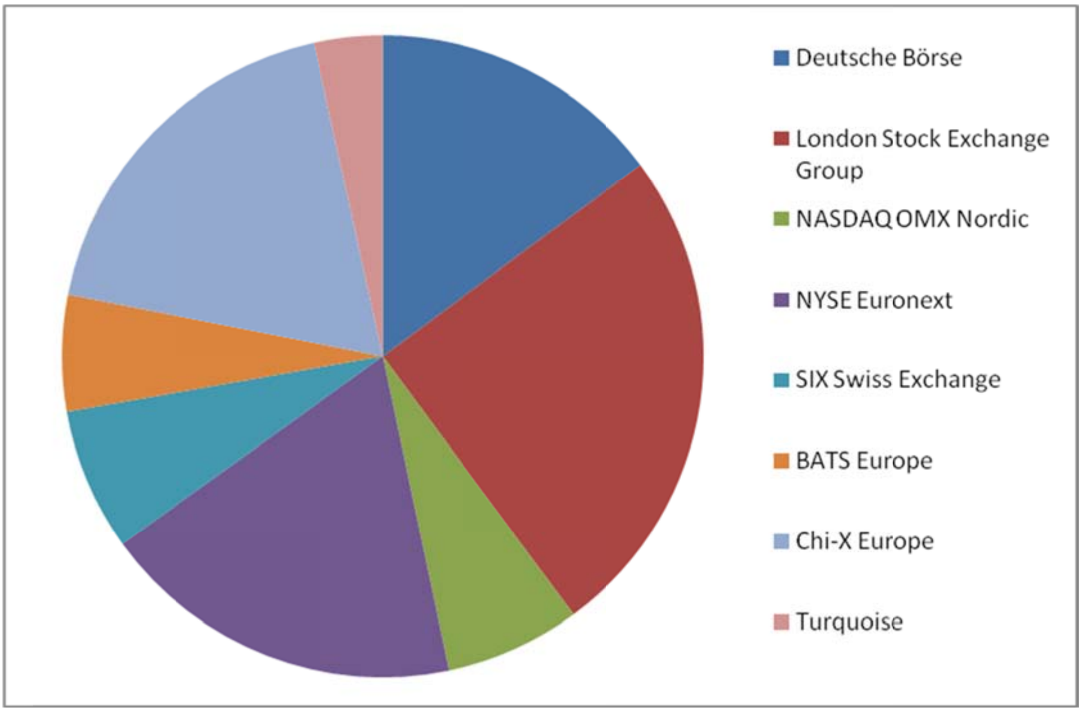

In addition to regulatory support, the emergence of alternative venues was facilitated by technological advances. This explains the introduction of electronic order-matching systems since the 1990's. Though their turnover was initially much less than on traditional exchanges, computer technology meant they could offer a lower cost, order-matching system by cutting out the intermediaries, such as broker-dealers.

The development and coexistence of several trading platforms has led to market fragmentation, i.e., quotes and depth for one stock are dispersed among market venues. In response to such fragmentation, market participants need to monitor prices and volumes as rapidly as possible. Thus, they search for the most attractive bids and offers across venues, splitting orders to reduce price impact and placing market and limit orders where dealing spread and market depth conditions are favorable. To seize trading opportunities before the market has moved away, speed is the secret of success. High frequency trading algorithms thus prove instrumental in helping traders coping with market fragmentation.

II.3) HFT algorithms can mitigate traders cognition limits

Even if trading is centralized, traders still have to monitor many sources of information: spreads and orders, the time series of trades, news flow, the position of their own book and the risk exposure of their institution, etc. Manually analyzing the various sources of information could delay their trades, which could in turn hinder the realization of gains from trades. In this context, trading algorithms can compensate the limited rationality of market participants and thus improve market liquidity. Biais, Hombert and Weill (2010) offer a theoretical analysis of these effects. They model the situation where the market is hit by an aggregate liquidity shock,in which a large fraction of the investors' population is subject to a transient decline in the willingness to hold the asset. Limits to traders' cognition make it difficult to reallocate the asset efficiently to the traders valuing it the most. Trading algorithms help mitigate this market imperfection. Algorithmic traders buy early after the shock, when prices are relatively low, while simultaneously placing limit orders to sell intended to execute later when the price will have reverted upward. Such trading strategies, which offer inter-temporal intermediation, are akin to a form of market making,similar to that analyzed by Grossman Miller (1988)). Biais, Hombert and Weill (2010) show that the liquidity supply strategies of the algo-traders involve frequent return trips, in which one algo-trader will buy and then resell to another algo-trader, and so on. Biais, Hombert and Weill (2010) also show that, when implementing their liquidity supply strategies, algorithmic traders will initially place a sequence of buy orders, at higher and higher prices. This leads to patterns tantamount to order splitting and short term momentum trading.

These theoretical results are in line with empirical findings. Brogaard (2010) finds that algorithmic traders tend not to withdraw from the market after shocks. Using a methodology inspired by Biais, Hillion and Spatt (1995), Hendershott and Riordan (2010) find that algorithms tend to provide liquidity when it is scarce and rewarded. Kirilenko et al (2010) find that HFTs "seem to be providing liquidity by putting resting orders in the anticipated direction of price moves." Note also that, consistent with the empirical results of Brogaard (2010) and Kirilenko et al (2010), algorithmic traders in Biais, Hombert and Weill (2010) use both limit and market orders. For example, Brogaard (2010) finds that, in his Nasdaq sample, in 50.4% of the trades the liquidity taker (i.e., the trader placing the market order) was a high-frequency trader, and that in 51.4% of the trades the liquidity supplier (i.e., the trader placing the limit order) was a high-frequency trader.

III) The Darker Arts

III.1) Manipulation

Some of the tactics used by high frequency traders are designed to mask deals and prevent other market participants from discovering and exploiting their trading intentions. This is also the motivation for the establishment of "Dark Pools". But the line between defense and offence is blurred.To the extent that a relatively small number of HF traders account for a significant fraction of turnover in a market, they may also engage in forms of market manipulation. Three such forms are "stuffing", "smoking" and "spoofing".

"Stuffing" involves HF traders submitting an unwieldly number of orders to the market. This generates congestion. In these conditions, access to the market for slow traders (non-HF traders) is impaired. They do not have a clear view of the current status of trading and it is difficult for them to execute trades. Meanwhile, fast traders who better understand what is going on and have superior access to the market engine, are able to execute profitable trades at the slow traders' expense.

When engaging in "smoking", high frequency traders first post alluring limit orders to attract slow traders. Then they rapidly revise these orders onto less generous terms, hoping to execute profitably against the incoming flow of slow traders' market orders.

Yet another strategy has been nicknamed "spoofing." Suppose the high frequency trader's true intention is to buy. Paradoxically, he or she will initially place limit orders to sell in the order book. These orders are not intended to be executed. Therefore they are placed above the best ask. And, since the high frequency trader is faster than the other market participants, he or she can rest assured he or she will have time to cancel the sell orders before they are executed if good news reach the market. With this assurance in mind, the high frequency trader places a sequence of limit sell orders above the best ask, potentially for very large amounts. The hope is to scare the market and induce some naïve participant to sell … against the limit order to buy the high frequency trader will have discretely placed meanwhile.

III.2) Adverse selection

Advantage is conferred on any market participant having early and privileged access to relevant information and the ability to trade on it before others. High Frequency traders are therefore prepared to invest in costly equipment and incur co-location charges set by exchanges to gain this advantage.

Thanks to these investments, HFTs orders rapidly impound all available information. Thus, Brogaard (2010) and Hendershott and Riordan (2010) find that algorithmic trades lead price discovery. For example, using the Hasbrouck (1995) variance decomposition methodology, Hendershott and Riordan (2010) find the contribution of HFT to price discovery to be greater than its contribution to trading volume. Correspondingly, HFT market orders have a greater permanent price impact than "slow human" trades.[12] For example, using the vector-auto-regression approach of Hasbrouck (1001a & b), Hendershott and Riordan (2010) find that HFT trades have a more than 20% larger permanent price impact than human trades. Figure 3, borrowed from Hendershott and Riordan (2009), illustrates this point by plotting the cumulative impulse response function (measuring the informational impact of trades) for HFT and human trades.

In other words, market orders from HFTs impound more information that human orders. Indeed, Kirilenko et al (2010) find empirically that "possibly due to their speed advantage or superior ability to predict price changes, high frequency traders are able to buy right as the prices are about to increase." The flip side of this informational edge is the corresponding adverse selection cost borne by "slow human" traders. [13]

First consider the case where market orders can be placed by "slow humans" or "fast computers" and hit quotes placed by human traders. Biais, Foucault &Moinas (2010) offer a theoretical analysis of this situation. To capture some of the main characteristics of trading algorithms, they assume that "fast" traders i) find trading opportunities more often than slow traders (consistent with our discussion of the motivations for HFT in the previous section), and also ii) observe information relevant to value stocks faster than slow traders. Because of i), HFT algorithms can increase the gains from trade realized in the market. But because of ii), they worsen adverse selection. This leads to greater realized spreads for slow traders, which

[12] This contrasts with the result by Hasbrouck (1998) that program trades on the NYSE in 1990 and 1991 had similar price impacts as human trades. Note that, while HFT orders are submitted faster than human orders, there was no such difference in latency between program trades and human trades in 1990 & 1991.
[13] The seminal studies of adverse selection costs in a microstructure context are Glosten and Milgrom (1985) and Kyle (1985).

reduces their gains from trade. In this context, algorithmic trading can be understood as a negative externality imposed on slow traders. When this externality is strong enough, it reduces their market participation and can even generate a market breakdown for slow traders. This might be the rationale for the empirical findings of Jovanovic and Menkveld (2010) that the entry of an HFT on the market for Dutch stocks led to a 13% decline in trading volume.

Figure 3:

Cumulative impulse response function (measuring the informational impact of trades) for HFT and human trades. Hendershott and Riordan (2009).

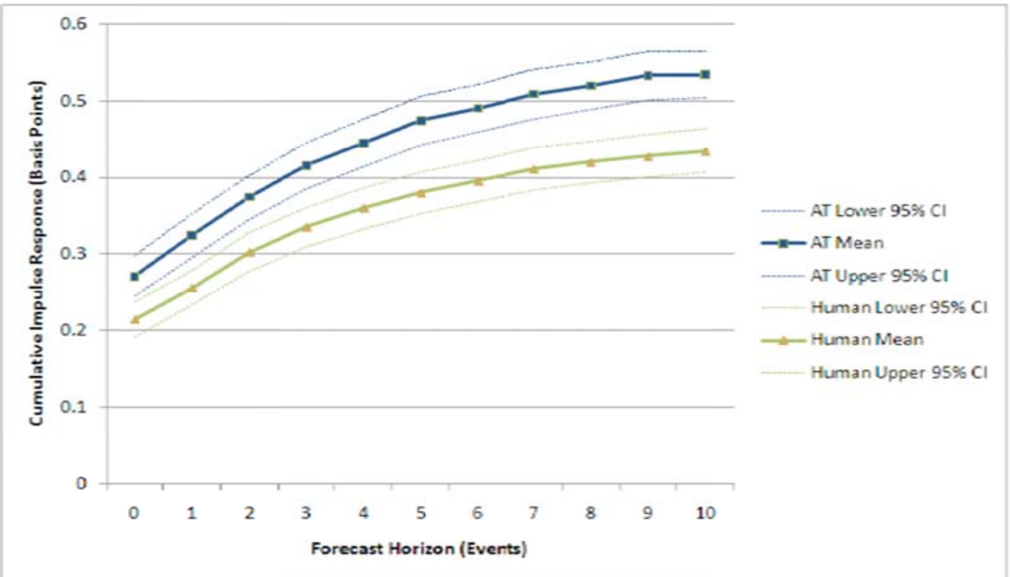

Now turn to the case where HFTs post market orders as well as limit orders. In this context again, if high frequency traders are more informed (or have access to information earlier) than human traders, then the former impose adverse selection costs on the latter. To illustrate this, consider the following example. The current best ask quote for stock XYZ is EUR 100. It has been posted by a "slow human" trader. Suddenly, public news hit the market. They raise the fundamental value of the stock to EUR 100.1. Algorithmic traders react extremely fast to these news, and hit the ask quote in stock XYZ, before the slow trader had the time to cancel his quote. Such algorithmic trades will indeed lead price discovery, but for the slow trader who placed an ask quote at 100 they create an adverse selection problem.

Consider a variant of the above case where the current best ask quote at 100, was posted by an algo-trader. The next best quote, at 100.01, was posted by a “slow human” trader. As long as no news hit the market, the HFT leaves its limit order to sell in the book and, when there is a market order to buy, is executed. But, as soon as good news hit the market, the algo-trader cancels its limit order to sell, and the “slow human” limit order to sell is executed.
This, again, generates an adverse selection problem for the human limit order trader.

The empirical results of Chaboud, Chiquoine, Hjalmarsson and vega (2009) are consistent with these analyses. They find that the permanent price impact of market orders hitting quotes placed by humans is greater than that of market orders hitting quotes placed by computers. That is, while for human traders, limit order executions are (to some extent) bad news, for computers limit order executions are profitable.

III.3) Imperfect Competition

As mentioned above, acquiring the technology necessary for HFT involves significant fixed costs. Thus, while for investors completing a large volume of trading, investing in this technology is profitable, it is not for less active investors. Biais, Foucault &Moinas (2010) offer a theoretical analysis of equilibrium investment in such a context. They show that in equilibrium there is a non-level playing field wherea small number of very actively trading fast players will coexist with slower traders conducting less many trades.[14]

This is consistent with the results of a recent survey (by Aite group) documenting that in the US equity market high-frequency trading firms represent 2% of the approximately 20,000 firms operating today, but account for 73% of all equity trading volume.Similarly,Jovanovic and Menkveld (2010) find that for the stocks of the Dutch index, one high frequency trader participated in more than 35% of the trades on Chi-X.

These remarks suggest that algorithmic trading could reduce the extent to which liquidity supply is competitive. Also in support of this claim, note that the adverse selection problem generated by algorithms for slow traders (and illustrated above in the example of stock XYZ) might deter them traders from posting limit orders.

III.4) Systemic Risk

To the extent that high frequency trades would be highly correlated and rely on similar strategies, they might contribute to destabilizing markets.[15] Now, in their empirical study of algorithmic trading in the foreign exchange market in 2006/7, Chaboud et al (2009) find that algorithmic trades tend to be correlated, suggesting that the algorithmic strategies used in the market are not as diverse as those used by non-algorithmic traders. To illustrate their analysis, they discuss an interesting case study of the dynamics of the Dollar Yen market on August 16, 2007. As discussed in their paper, and illustrated in Figure 4 (which we borrow from their study), on that day the Japanese yen appreciated sharply against the US dollar between 6:00 a.m and 12:00 p.m. NY time. The sharp exchange rate movements happened when computers, as a group, aggressively sold dollars and bought yen. Chaboud et al (2009) find that computers, during this episode, hit human quotes much more than computer quotes, and the market orders generated by computers were far more correlated than those generated by humans. The authors also find that, after 12:00 p.m. human traders began to buy dollars and the appreciation of the yen against the dollar was partially reversed. This narrative is consistent with the view that correlated order flow from computers destabilized the dollar on that day, but that this price impact was later reversed by human order flow.

---

[14] This arises even if co-location facilities are offered on a non-discriminatory basis, i.e., at the same price to all market participants, as long as the price of co-location or the other investments required by HFT (hardware, people, connexions, …) have a fixed cost component.

[15] The August 2007 mini crash offers a good illustration of how correlated strategies can generate systemic risk: At that time, many quants were using similar strategies. Thus, they were simultaneously hit by a shock, and reacted similarly, which generated a downward spiral in the market.

Figure 4: Pricing dynamics and order flow in the Dollar yen market on August 16, 2007. Chaboud et al (2009).

Figure 4, Panel A: Computer taker order flow

Figure 4, Panel B: Human taker order flow

Also suggestive of correlation between HFT strategies, using the transition matrix methodology developed in Biais, Hillion and Spatt (1995), Brogaard(2010) finds that the serial autocorrelation in order types is more pronounced among high frequency traders than “humans”.

Thus, it is possible that algorithmic traders, using correlated trading strategies, would participate in a large fraction of the trades and impact the market price. In this context, shocks hitting key algorithmic traders might affect the entire market. Note further that, as discussed above, slow human traders are exposed to adverse selection when dealing with high frequency traders. Such adverse selection might deter slow human traders from entering the market to provide liquidity to HFTs when the latter would be subject to a shock. Thus, algorithmic trading could create the scope for systemic risk.

What shocks might trigger such events? It could be shocks related to the risky positions of algorithmic trading firms. Khandani and Lo (2008) analyze the mini-crash that took place on the NYSE during the week of August 6, 2007. Hedge funds, learning they had incurred severe losses on their real estate related positions, underwent a decline in their ability to hold risky positions. To cope with such financial stress, they sold stocks on the NYSE. Because several hedge funds were hit in a similar fashion, there was a wave of sales. This depressed the prices, which triggered more sales, resulting in a severe (but short-lived) decline in the index.

At the time, algorithmic trading was not as developed as it is now, and it played very little role in the 2007 minicrash. But, reasoning by analogy, one cannot rule out that a comparable crisis could hit algorithmic trading firms. Note that the likelihood of such problems would be particularly high if algorithmic trading firms had a small capital base. Now, unlike banks, pure-play algorithmic trading firms are only very mildly regulated and face no capital requirements.

The trigger could alternatively be technical problems, regarding computers or connections to the market, a form of operational risk. One could also interpret the flash-crash of May 6, 2010, as a manifestation of the systemic risk induced by algorithmic trading.

As documented in the CFTC-SEC (2010) report, due to worries about sovereign risk, the depth in the book that morning was limited. Yet a large fundamental trader decided to execute a large sale program with a trading algorithm. Unfortunately, this algorithm was not programmed to reduce the scale of its trade in case of large price impact. Its orders were initially absorbed by high frequency traders and other investors. But, after this initial phase, high frequency traders reversed their positions. This increased the selling pressure and increased the drop in prices. This chain of events is illustrated in the last panel of Figure 1.B, borrowed from Kirilenko et al (2010). The Dow index dropped by almost 1,000 points in a matter of minutes. It is only after trading was halted by market organizers, and participants had the time to check that no negative fundamental information could justify such a drop in prices, that the market recovered. In the terminology of the systemic risk scenario we proposed above, the initial algorithmic selling program can be interpreted as the shock that hit the high frequency traders and triggered the systemic risk event. Our reading of the events is consistent with Kirilenko et al (2010) who conclude that “High Frequency Traders did not trigger the Flash-Crash, but their response to the unusually large selling pressure on that day exacerbated market volatility.” Note also that, while the flash-crash was not initially triggered by high-frequency traders, it was triggered by an algorithm, and, during the few minutes that the flash-crash lasted, a large fraction of the order flow stemmed from HFT firms.

IV) Overall assessment

IV.1) Informational efficiency

To the extent that it brings prices of related assets in line, HFT improves informational efficiency. Also, to the extent that it enables traders to better and faster process relevant information, HFT enhances price discovery. Several empirical studies offer evidence consistent with these claims. This leaves the question open whether HFT improves informational efficiency in a way that is socially useful. If it only results in information getting impounded in prices a few seconds or minutes earlier than it would be without HFT, then maybe it does not contribute significantly to making economic decisions more efficient.

IV.2) Liquidity

Empirically, HFT seems to be associated with elevated trading volume. Does this imply that it increases liquidity? Let's define liquidity as the ability to rapidly conduct transactions that realize mutual gains from trade. Liquidity, thus defined, is not always equivalent to volume.

In the theoretical analysis of Biais, Hombert& Weill (2010), trading volume in the perfect market (where all gains from trade are immediately exploited) is lower than in the imperfect market (where, because of cognition limits, it takes more time to arrange mutually beneficial trades.) Large trading volume in the imperfect market is generated by the return trips of the algo-traders. These trades are conducted to help the market cope with the imperfection, but don't fully mitigate it.

In Biais, Foucault and Moinas (2011), HFT has ambiguous consequences, since on one hand it can help traders finding counterparties and realizing gains from trades, but on the other hand it can generate adverse selection problems, preventing the realization of gains from trade. In this context, the level of algorithmic trading that maximizes trading volume is not identical to the level that maximizes utilitarian welfare.

Now consider a dynamic market under asymmetric information, where computers and humans can place limit and market orderS. Suppose, at time t, a human investor, with a low private value, contacts the market to sell 100 shares. Suppose also that, at time $t + s$, another investor will arrive and seek to buy 100 shares. Finally suppose that the initial seller cannot remain on the market and continuously monitor the arrival of orders and the changes in fundamental value. In this context, if there is no HFT, and if information asymmetry is not too severe, the initial seller will leave a limit order to sell for 100 shares in the book. At time $t+s$, the buyer will hit this limit order and buy. Total trading volume will be equal to 100 shares. Now turn to the case where HFT market makers operate in this market. When the initial seller contacts the market, the HFT market maker will buy from him, and then sell back to the buyer at $t+s$. In this context, trading volume will be larger than the 100 shares traded without HFT, and yet gains from trade, which result from the exchange of shares between the initial seller and the final buyer, will not be increased. This reasoning is consistent with the theoretical analysis of Cartea and Penalva (2010), where HFT increases volume by intermediating trades. Note further that, if the HFT earns a margin (a bid-ask spread) on the trade, its intervention might even reduce gains from trade, if the seller reacts to the transactions cost by lowering the amount traded, say to 95 shares instead of 100 – yet, even then, total volume is larger with HFT than without. One might wonder whether the initial seller could attempt to avoid paying the spread, by rejecting the offer of the HFT and placing his own limit order to sell. This strategy might turn out to be unprofitable if the continuous presence of the HFT, which enables it to observe changes in fundamental value and quickly react, generates adverse selection cost for the slow trader placing a limit order to sell.[16]

V) Perspectives and policy implications

V.1) Likely evolution under laissez faire

If regulators and policy makers choose not intervene, we expect one of the following scenarii to unfold.

It may turn out that algorithmic trading will not trigger any major systemic crisis.[17] In that case, HFT is likely to continue to thrive. Banks, hedge funds and "pure play" HFT firms will continue to engage in an arms' race. Efforts to minimize latency will ultimately find their limits, since the latency is bounded below by zero. But HFT firms will also continue to develop increasingly sophisticated and rapid trading algorithms. These evolutions may, in the end, benefit markets and investors as a whole, by improving price discovery and liquidity. But we think there is a significant risk that they could, on balance, be detrimental to slow traders, in particular due to adverse selection. In that case, final investors and their investment managers will need to seek protection from HFT. To do so, buy-side firms will continue to develop their own trading algorithms, and thus participate in the arms' race.[18] Alternatively, slow investors will retreat from electronic markets where they are too exposed to HFT. They will increasingly rely on dark pools and on OTC crossing, where their orders are hidden and therefore not exposed to predation by HFT. Such an evolution, itself, would not be without potential costs: Order flow diversion from transparent exchanges could hinder the informational aggregation function of markets, and internalization of order flow execution inside banks could raise agency issues.

On the other hand, it is possible that there will be a crisis related to HFT. As discussed above, this could be triggered by a form of operational risk (hardware failure or dysfunctional algorithms). It could alternatively be triggered by an outside shock hitting HFT firms. These firms would then try to rapidly close their positions. Similarly to what happened during the mini-crash of August 2007, simultaneous attempts to do so, stemming from high frequency traders following correlated strategies, could lead to a downward price spiral (Gromb and Vayanos (2002) show how institutional constraints on arbitrageurs can amplify shocks). In a very short span of time, HFT could loose millions of dollars, as happened during the flash-crash of May 2010. This could push lightly capitalized HFT firms into bankruptcy and generate defaults. Since HFT firms intensively trade with one another, defaults could propagate due to counterparty problems. The latter would be all the more difficult to solve that HFT firms trade in multiple markets, involving different clearing and settlement systems.

To steer clear from the risks identified in these two scenarii, it would be prudent to put in place adequate regulations. As written by Kirilenko et al (2010) in their conclusion:

---

[16] See our analysis of such adverse selection costs in Subsection III.2) above.

[17] This would be consistent with the observation that, so far, any flash-crash or mini-crash has been rapidly stabilized and followed by price reversion.

[18] As with Lewis Caroll's Red Queen, since all the others are moving fast, to maintain your current position you also have to make rapid progress.

"technological innovation is critical for market development. However, as markets change, appropriate safeguards must be implemented to keep pace with trading practices enabled by advances in technology."

V.2) Regulating HFT

Surveillance and taxes

HFT can generate i) mutually beneficial gains from trade as well as ii) adverse selection. While both can lead to profits for HFT, only the former are socially useful, while the latter induce zero-sum trading games. As shown by Biais, Foucault and Moinas (2010), when deciding whether to invest in HFT infrastructures financial firms take into account both sources of profitability. This contrasts with the socially optimal level of investment in HFT infrastructure, which would only be increasing in the profits from i), not those from ii). Therefore, in equilibrium, there will be over-investment in HFT infrastructures.[19]

An appropriate policy response could be to tax HFT, for example by taxing collocation. Such taxes would be akin to Pigovian taxes, leading high frequency traders to internalize the adverse selection costs they impose on slow traders. Such taxation, however, could be difficult to implement. One of the problems in practice would be the competition between differently regulated areas, another would be the uncertainty about what the level of the tax should be.

Yet, from a pragmatic perspective, taxes on HFT could also be justified and implemented as a way to fund market monitoring – thus killing two birds with the same stone. HFT generates huge order flows, making it very difficult and costly for regulators and market organizers to monitor trading. At the same time, HFT potentially creates systemic risk, which increases the social value of market surveillance and regulatory monitoring. It is therefore important to ensure that regulators and market organizers have the means (in terms of staff, connections and computer equipment) to complete their surveillance task. The level of resources needed to fund this public good, would be a good benchmark to set the target level of proceeds from taxes on HFT. This public policy could be designed and implemented in collaboration between the European Union (e.g., in the context of the revision of MIFID) and the US regulatory authority (SEC & FINRA).

Slowing the market

Another way to deter wasting resources on socially useless investments, would be to impose a minimum latency, e.g., one tenth of a second. Some traders oppose minimum latency requirements on the grounds that such limits are backward and hopeless attempts to avoid technological progress. The same criticism applies to speed limits on the roads. And it is hard to believe that going from a latency of a millisecond to a latency of one tenth of a second would significantly hinder the information aggregation function of the market. It is in fact possible that allowing infinitesimal latency reduces gains from trade.
Vayanos (1999) analyzes the dynamic strategies of informed agents as a function of the length of the time interval between trades (h). To reduce price impact, strategic agents split their

[19] This is consistent with the view of investment in HFT technology as an arm's race, in which the Nash equilibrium yields investment above its socially optimal level. QUOTE RICK GREEN S PAPER

trades. Since this reduces their ability to trade out of their endowment shocks, this reduces the gains from trade achieved in the marketplace. Vayanos (1999) shows that the welfare loss is maximized when *h* goes to 0. The intuition is the following. If the time between trades was longer, a market maker could accommodate orders without fearing that they would immediately be followed by additional trades altering further the value of the stock.

A more drastic way of slowing the market is to trigger trading halts. As illustrated by the flash-crash of May 6, 2010, HFT could increase the risk of sharp market movements, unwarranted by fundamentals, and giving rise to trades at disequilibrium prices. Circuit-breakers, which halt trading when price changes are unusually large, could mitigate such mispricing. They could give traders the time it takes to analyze the value of the asset, and help the market revert to equilibrium. In particular they can be necessary to leave time for human agents to analyze the market and intervene to offset mispricing induced by dysfunctional computer programs.

But, for such safeguards to be efficient, they have to apply across trading platforms and related securities, as opposed to being stock and exchange specific. As illustrated also by the flash-crash of May 6, 2010, since HFT algorithms arbitrage across markets, halting trades in the underlying while continuing it in the derivative (or vice versa) can be very damaging.

The SEC, the Finra and the US exchanges have implemented market wide circuit breakers for the stocks in the 2010 Russell 1000 Index and to a list of exchange-traded products, including those that track broad-based stock indexes. These market rules were implemented on a pilot basis. We think they should be upheld in the long term.

Regulatory oversight and capital requirements

The European Commission has included the analysis of HFTs in its review of the Market In Financial Instruments Directive. It considers the possibility to subject HFT firms to regulatory oversight and capital requirements. This would help prevent systemic risk creation by HFT firms. First, capital buffers would reduce the likelihood that HFT firms would be destabilized by liquidity shocks and would in turn destabilize their counterparties. Second, capital requirements could increase the "skin in the game" of the manager owners of HFT firms, and reduce the moral hazard problem associated with limited liability.

Similarly to the case of bank, setting optimal capital requirements for HFT firms will require a better understanding of their role in the market, and their potential systemic consequences. In particular, stress tests should be conducted, to evaluate the likely consequence of operational risk events, the failure of some HFT firms, and the potential for contagion via counterparty defaults.

The necessary effort to better understand the nature and consequences of HFT should also be exerted by academics. While, some very informative empirical studies are already available, it is often difficult to draw causal inferences because the variables of interest (HFT activity, market liquidity, trading volume, adverse selection,…) are jointly endogenous. Hendershott, Jones and Menkveld (2010) take an intriguing instrumental variable approach to deal with the endogeneity problem. More econometric work is needed to study these issues further.

References

Biais, B., T. Foucault & S. Moinas (2010), “Equilibrium algorithmic trading,” Working paper, Toulouse School of Economics (IDEI).

Biais, B., J. Hombert, P.O. Weill (2010), “Trading and liquidity with limited cognition,” Working paper, Toulouse School of Economics (IDEI).

Biais, B., C. Bisière and C. Spatt, 2010, “Imperfect Competition in Financial Markets: An Empirical Study of Island and Nasdaq”, Management Science, 2237–2250.

Biais, B., P. Hillion and C..Spatt (1995), “An empirical analysis of the order book and order flow on the Paris Bourse,” *Journal of Finance*, 1655-1689.

Brogaard, J. (2010), “High frequency trading and its impact on market quality,” Working paper, Northwestern University.

CFTC-SEC (2010), “Findings regarding the market events of May 6, 2010.”

Cartea, A. and J. Penalva (2010), “Where is the Value in High Frequency Trading?”, Working paper, Universidad Carlos III, Madrid.

Chaboud, A., E. Hjalmarsson, C. Vega & B. Chiquoine (2009), “Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market”, Federal Reserve Board International Finance Discussion Paper No. 980

Chlistalla, M. (2011), “High frequency trading?Better than its reputation?”Deutsche Bank Research Briefing, February 7, 2011.

Glosten, L. and P. Milgrom (1985), “Bid, ask, and transactions prices in a specialist market with heterogeneously informed agents.” *Journal of Financial Economics*, 71-100.

Gromb., D. & D. Vayanos (2002) “Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66:361-407.

Grossman, S. and M. Miller (1988), “Liquidity and market structure,” *Journal of Finance*.

Harris, L., .Sofianos& J. Shapiro (1994), “Program trading and intraday volatility,” *The Review of Financial Studies*, 653-685.

Hasbrouck, J. (1991a), “Measuring the information content of trades,”*Journal of Finance*,179-207.

Hasbrouck, J. (1991b), “The summary informativeness of stock trades: An econometric analysis,”*Review of Financial Studies*,571-595.

Hasbrouck, J. (1995), “One security, many markets : Determining the contributions to price discovery,”*Journal of Finance*,1175-1199.

Hasbrouck, J.(1996), “Order characteristics and stock price evolution: An application to program trading.” *Journal of Financial Economics*, 129-149.

Hasbrouck, J. and G. Saar (2011), “Technology and Liquidity Provision: The Blurring of Traditional Definitions,” forthcoming *Journal of Financial Markets*.

Hendershott, T. & R. Riordan (2009), “Algorithmic trading and information,” Working paper, Net institute, Working paper # 09-08.

Hendershott, T., C. Jones and A. Menkveld (2010) “Does Algorithmic Trading Improve Liquidity?”,*Journal of Finance*.

Jovanovic, B., & A. Menkveld (2010), “Middlemen in limit order markets,” Working paper, New York University.

Khandani, A. & A. Lo (2008), “What happened to the quants in August 2007?”, forthcoming *Journal of Financial markets*.

Kirilenko, A., A.S. Kyle, M. Samadi& T. Tuzun (2010), “The impact of high-frequency trading on an electronic market,” Working paper, University of Maryland.

Kyle, A. (1985), “Continuous auction and insider trading;” *Econometrica*, 1315-1335.

Menkveld, A. (2011), “High-Frequency-Trading and the New-Market Makers,” Working Paper,, VU University, Amsterdam.

Pagano, M. (1989a), “Endogenous Market Thinness and Stock Price Volatility,” *The Review of Economic Studies*, Vol. 56, 2, 269-287.

Pagano, M. (1989b), “Trading Volume and Asset Liquidity,” *The Quarterly Journal of Economics*, Vol.CIV, 2.

Schapiro., M. (2010) "Remarks Before the Security Traders Association". http://www.sec.gov/news/speech/2010/spch092210mls.htm.

Vayanos, D., (1999) “Strategic Trading and Welfare in a Dynamic Market,” *Review of Economic Studies*, 66, 219-254.