The Good, the Bad, and the Ugly of Automated High-Frequency Trading

Tommi A. Vuorenmaa[1]

Valo Research and Trading

June, 2012

Abstract

This article discusses the pros and cons of automated high-frequency trading (HFT). There appears to be much confusion of whether HFT is "good, bad, or ugly." In the terminology of this article, the "Ugly" category consists mostly of popularized negative writing against HFT appearing in media. The category we label as "Bad" consists of more detailed research arguments against HFT. Perhaps surprisingly to non-professionals, the "Good" arguments outweigh the others by a milestone in academic studies. These arguments are often ignored in media and should be brought to fore of the discussion to be fair. We review the commonly presented arguments as neutrally as possible and attempt to bring some additional insight to the discussion.

INTRODUCTION

This article discusses the pros and cons of automated high-frequency trading (HFT). The Securities and Exchange Commission (SEC) calls HFT "One of the most significant market structure developments in recent years" [see SEC (2010)]. The main motivation for writing this article is the heated debate over the positive and negative effects of HFT on the quality of financial markets.[2] By quality, we mean primarily liquidity, price efficiency, and price discovery – overall, market health. In

1 E-mail: tommi.vuorenmaa@valotrading.com

Comments are welcome. I thank Aki Saarinen at Valo for constructive comments on the written language.

2 Wall Street Journal. (October, 2011). "Does High-Speed Trading Hurt the Small Investor?"

some cases, the debate is getting blown out of proportions, lacking good sense.[3] In recent years, more and more articles studying this topic have been written, but there still appears to be much confusion of whether HFT is “Good, Bad, or Ugly.” On the one hand, for example Gomber et al. (2011) highlight the beneficial aspects of HFT and see the perceived problems to be a largely a result of the US market structure. On the other hand, the SEC (and the regulatory authorities in Europe) is scrutinizing HFT in fears it puts “less tech-savvy investors at a disadvantage.”[4] Behind this fear lies the perhaps unconscious worry that computers are about to run over humans in investing as they have done in other demanding fields and complex games, such as chess.[5] [6] [7]

We review the current literature on HFT and the closely related algorithmic trading (AT), but we do not claim to be comprehensive. We emphasize academic studies, because their views of this topic seem to be the most neutral. Importantly, we do not merely present the evidence but also discuss it from our own point of view. We hope this gives fresh perspective and brings research in HFT and related fields to the fore.

Defining HFT

Before discussing the pros and cons of HFT, one thing should be cleared right away: HFT and AT are not the same concepts. Many academic authors consider HFT as a subset of AT [see, e.g., Hendershott and Riordan (2011)], but this practice should be challenged more. We see HFT as a subset of automated (not algorithmic) trading, a view which we share for example with Markets Committee (2011).[8] Automated trading, in our view, includes both HFT and AT as its subsets – both require high level of automation, but for different purposes, as explained below. Different definitions of

3 YouTube: Joe Saluzzi of Themis Trading and Irene Aldridge of ABLE Alpha Trading discuss HFT on CNBC.
4 Financial Times. (July, 2009). "High-Frequency Trading Under Scrutiny."
5 Wired Magazine. (December, 2010). "Algorithms Take Control of Wall Street."
6 BBC. (September, 2011). "Quant Trading: How Mathematicians Rule the Markets."
7 Gasparov, Garry. (February, 2010). "The Chess Master and the Computer." New York Review of Books.
8 To emphasis the view we take, in the title of this article we use “automated HFT” instead of simply “HFT.” Some authors use automated and algorithmic interchangeably [see, e.g., Zhang and Powell (2011)], which is confusing. Some proprietary trading firms propose to use the term “automated professional traders” [see Connell et al. (2010)].

AT and HFT used in the literature are collected in Gomber et al. (2011). Ignoring the conceptual differences have probably added to the confusion surrounding HFT. AT and HFT being different entities, their pros and cons are not necessarily the same.

Traditionally, an automated trading model (HFT or slower) is about determining whether a trade should be placed, while an algorithmic trading (AT) model is about determining how to place a trade in order to minimize trading costs. More explicitly, automated HFT is about profit-maximizing using trading strategies specifically developed for this task, while AT is about minimizing transaction costs of trading large blocks of assets.[9] A number of authors, for example SEC (2010) and Brogaard (2012), list specific features required of HFT.[10] Simply put, HFT is "fast" automated trading. How fast the trading must be depends on the trading strategy. Typically, HFT involves holding times of assets of only minutes or seconds, the technologically most savvy trading firms working in the sub-second range (milliseconds, or even faster).

Scope of this article

Donefer (2010) divides the users of automated trading strategies into four groups: (i) liquidity seekers; (ii) automated market makers; (iii) statistical arbitrage strategy traders, some of whom are gamers or predators; and (iv) rebate seekers. In this article, we accept a rather broad definition of HFT: automated trading executed at intra-daily intervals, but such that it excludes trading for the purposes of minimizing transaction costs. Our definition of HFT thus includes the three last groups and excludes the first.

9 Bertsimas and Lo (1998) study the dynamic optimization problem of how large positions of assets should be accumulated or liquidated through breaking them into smaller pieces to minimize the expected cost of execution.

10 "[HFT] typically is used to refer to professional traders acting in a proprietary capacity that engage in strategies that generate a large number of trades on a daily basis. Other characteristics often attributed to proprietary firms engaged in HFT are: (1) the use of extraordinarily high-speed and sophisticated computer programs for generating, routing, and executing orders; (2) use of co-location services and individual data feeds offered by exchanges and others to minimize network and other types of latencies; (3) very short time-frames for establishing and liquidating positions; (4) the submission of numerous orders that are cancelled shortly after submission; and (5) ending the trading day in as close to a flat position as possible (that is, not carrying significant, unhedged positions over-night)." [SEC (2010)]

We focus mostly on the present and, also to some extent, the future of HFT. There exist several good articles describing the historical background on HFT. We refer the interested reader to McGowan (2010) for a good introduction to the "rise of the machines." Shorter descriptions can found found for example in Biais and Woolley (2011) and Vuorenmaa (2012a). Technological developments are discussed in Cliff (2011) and Cliff et al. (2011). These articles help to put HFT in the right context. Another limitation of our article is that we do not discuss the role of financial market fragmentation or different models of market microstructure in much detail. For the former, we refer the reader to O'Hara and Ye (2011) and Vuorenmaa (2012b), and for the latter to De Jong and Rindi (2009). These aspects are usefully intertwined in the theoretical framework of Pagnotta and Philippon (2012) who study the joint evolution of trading regulations, market fragmentation, and speed, and their effect on welfare.

THE UGLY

We now go directly to the core matter of the current article: reviewing and discussing the empirical evidence of HFT, and in particular whether the pros outweigh the cons. In our context, the category the "Ugly" includes perhaps the most serious objections raised against HFT. They are certainly the most popularized ones by media.[11] As they also tend to be the most concrete and feared objections, we begin by discussing them.

Flash crashes

The most well known, and most likely the "ugliest," argument against HFT is the Flash Crash of May 6, 2010. US stock market indices, stock-index futures, options, and exchange-traded funds then experienced a sudden and unprecedented increase in volatility. A drop of more than five percent in prices was followed by a rapid rebound – in about 20 minutes – of almost the same magnitude. Although an intraday event,

11 YouTube: Senator Kaufman discusses HFT on MSNBC's Dylan Ratigan Show.

Zigrand et al. (2011), among others, argue that it eroded confidence in markets for a considerable time to come, and was thus not inconsequential. A detailed description on what is thought to have taken place on May 6, 2010, can be found in the joint study of CFTC/SEC (2010), commissioned by Joint Advisory Committee on Emerging Regulatory Issues in the US.[12] Shorter popular accounts have appeared in MacKenzie (2011) and several finance industry magazines[13]. The importance of this event is emphasized by the fact that officially commissioned studies are rather rare – although by no means are fast crashes limited to the era of HFT, as discussed next.

Probably the most famous report was commissioned after the 1987 stock market crash. In that crash, the main blame was put on portfolio insurance strategies [see, e.g., Carlson (2006)]. Coupled with more standard stop-loss strategies, they created a feedback loop, bearing some resemblance to what happened on May 6, 2010. Markets Committee (2011) describe another (but less well known) event related to stop-losses that took place for Japanese Yen on March 17, 2011. Many HFT firms and market makers then reportedly withdrew from the market. The designated market makers did practically the same by submitting unrealistic quotes. Andrew Lo, chief investment strategist for AlphaSimplex Group, also finds parallels between the Flash Crash of May 6, 2010, to the so-called quant meltdown that took place in August, 2007, where "a rapid unwinding of a large equity-market-neutral portfolio may have caused losses in other portfolios using similar strategies" [see Khandani and Lo (2008)]. In 2007, however, automated trading was not yet in as a dominant position as it was in 2010.

Whom to blame for the Flash Crash of May 6, 2010, then?[14] Kirilenko et al. (2011) attempt to uncover this by studying this event using E-mini S&P 500 stock index future data. Their conclusion is the same that of CFTC/SEC (2010): HFT firms did not trigger the crash, and cannot be blamed for it, but their trading response to a

12 New York Times. (September, 2010). "Ex-Physicist Leads Flash Crash Inquiry."
13 Automated Trader Magazine. (2010). "What Just Happened?"
14 Buchanan, Mark. (May, 2012). "Flash-Crash Story Looks More Like a Fairy Tale."

sudden selling pressure by an institutional seller exacerbated volatility.[15] Markets Committee (2011) show evidence that algorithms were actively trading also in foreign exchange (FX) markets at the downward spiraling phase of the crash. HFT firms were aggressively trading in the direction of price changes, causing a “hot-potato” effect: a toxic asset changing hands from one to another in rapid succession. This could be due the HFT firms' trading style, which is different from the trading style of designated market makers (cf. the NYSE “specialists”). Because of liability, these designated market makers are typically less aggressive and accumulate large inventory positions. HFT firms, on the other hand, avoid taking such positions to minimize market risk.[16] Martinez and Roşu (2011) show in a theoretical framework that HFT firms can accomplish this by increasing their trading frequency. Thus, HFT firms may sometimes compete fiercely for liquidity and so amplify price volatility. This is what appears to have happened on May 6, 2010.[17] Ironically, HFT intraday risk minimization practices may actually increase systemic risk, as discussed later.

Empirical studies of the Flash Crash of May 6, 2010, present quite convincing evidence that HFT damaged liquidity and increased volatility on that particular day. The “doomsday dynamics” were apparently set in motion by a “run on a bank” type of self-fulling prophecy – a topic famously studied decades ago by sociologist Robert K. Merton [see Merton (1949)]. On May 6, 2010, the run seems to have started from an unwinding of a position reportedly executed by standard AT methods (see later). HFT firms reacted by increasing their trading, which signaled the algorithm to sell even more. This created a temporary “liquidity black hole,” evaporating liquidity, as discussed by Morris and Shin (2004). The effect was made worse by market makers avoiding to “catch a falling knife by its point.”[18] In the list of six different feedback

15 Not surprisingly, perhaps, because one of the authors works for CFTC.
16 Wall Street Journal. (October, 2011). "A Call to Pull Reins on Rapid-Fire Trade."
17 Interestingly, as noted in the ZeroHedge-website, Kirilenko et al. (2011) use a rather unusual definition of liquidity: in their paper, they argue that high-frequency traders supply liquidity by executing against the standing limit-orders. This is the opposite to the standard notion of liquidity supplying by placing limit-orders in the book. Whether this point of departure affects the conclusions regarding the role of the HFT firms in Flash Crash of May 6 is yet unclear.
18 Katz, Gary. (2011). "Relying on the 'Rules of the Game' in the U.S. Options Market."

loop systems described by Zigrand et al. (2011), the “volume feedback” system fits this description the best. More research on such feedback loops is warranted as we have limited experience on automated trading that is faster than manually possible.

Liquidity black holes are not an entirely new phenomena in high technology fields. Near catastrophic events, such as the Flash Crash of May 6, 2010, are what Perron (1984) categorize as a “normal accidents.” In Perron's view, such catastrophes are inevitable consequences of closely coupled processes that typically take place in “high-risk technologies.” Catastrophes are a latent feature of “socio-technological systems,” large-scale information technology systems supporting critical social and economic functions. They happen because risky events are often misinterpreted as being normal. The risk of catastrophes is typically not materialized until later by a process known as the “normalization of deviance” [see, e.g., Zigrand et al. (2011)]. According to Cliff (2010), the main problem of socio-technological systems is that failures are hard or impossible to predict. Failures could be controlled to some extent by simulation-based techniques, but these simulations must be realistic. Gsell (2008) uses simulations to demonstrate how AT impacts markets in terms of price formation and volatility. For such a framework to be useful in the sense just described, however, algorithms should not be considered in isolation, but in relation to other algorithms. Similarly, theoretical models like that of Carlin et al. (2005), which demonstrate how liquidity crises can arise from a breakdown of cooperation between different market participants, are typically equilibrium models and may thus not be complex enough. Farmer and Geanakoplos (2008), for example, discuss the limitations of such models.

In line with Perron's view of “normal accidents,” Johnson et al. (2012) find evidence of “mini flash crashes” that take place in the sub-second range, “at the limits of the monitoring capability of human response times.” By analyzing a set of 18,520 “ultrafast Black Swans,” Johnson et al. (2012) provide novel evidence of some risks involved in operating at the limits of, or beyond, human response times.[19] Although this topic warrants more thorough research, it signals that reasonable control should be exerted over trading at the market microstructure level where sub-seconds and cents matter. More exactly, the use of special orders and routing practices in the US markets may have played a key role in the perceived problems. Wood et al. (2012), in particular, suspect that the use of Intermarket Sweep Orders (ISOs) contributed to the occurrence of the Flash Crash of May 6, 2010. They find that the amount of ISOs was much higher on that day than normally and were mostly used by informed traders before the crash. Gomber et al. (2011) note that similar crashes have not taken place in Europe, where a different market microstructure is used. It would thus not make sense to regulate technology to get rid of the (mini) flash crashes. Instead, it seems that more efficiently working technology could, in many cases, save exchanges from experiencing message traffic overloading, for example. This should be kept in mind as authorities worldwide are facing pressure on regulating HFT.[20] [21]Simulating Black Swan events could offer a feasible way to assess risks involved, as we discuss below.

Digression: Black Swans and Dragon-Kings

As its name suggests, flash crashes happen fast. Another way to put it is to say that flash crashes are by their nature surprise phenomena. In the last decade, there has been a strong, renewed interest in extreme event modeling. A framework for it has been popularized especially by Nassim Nicholas Taleb. Building on much earlier work in academic literature, foremostly Mandelbrot (1963) and Fama (1965), Taleb (2007) calls surprise extreme events “Black Swans.” The original advocator of this framework, Mandelbrot, used often the term “wild” events, or the more descriptive “Noah-effect” (in reference to the biblical flood), in the context of extreme events.

19 Wired Magazine. (February, 2012). "Nanosecond Trading Could Make Markets Go Haywire."
20 Columbia Business Law Review. (2012). “Regulators to Enhance Scrutiny over High Frequency Trading Activities."
21 Connell, Liam, Richard B. Gorelick, Adam Nunes, and Cameron Smith. (September, 2010). "Market Safeguards."

Philosophical questions aside, the basic underlying theoretical assumption of Black Swans is that events are independent of each other. This means that they cannot be predicted. So, does the Flash Crash of May 6, 2010, qualify as a Black Swan in this sense? The answer is it all depends on what time-scale the event is being observed. In order to clarify this subtle point, assume for a moment that we sample some historical data covering May 6, 2010, at a daily frequency. As described above, the prices that experienced a sudden drop recovered after about 20 minutes of trading. Clearly, then, at a daily sampling frequency no Black Swan by this definition took place at all. Even at a hourly sampling frequency we would not have evidence of a Black Swan. But, if we would sample the data at a five minute frequency, say, we would see the evidence.

One could go further: sample at a second or, better yet, at a sub-second frequency, and the evidence starts to vanish again. The workings of the Flash Crash of May 6, 2010, was such that the changes in prices were not in fact independent of each other. There was a feedback loop that accumulated to a one large surprise event at some time-scale (in this case, 5 minutes, say). So where does this line of reasoning lead us? Clearly, Black Swans are more precisely speaking black or wild events only at certain time-scales. In other words, the blackness or wildness can be “tamed” by moving from a time-scale to another. This is actually a more general phenomena in financial markets: the changes of (logarithmic) prices tend to have probability distributions with heavy-tails when the time-scale is small, but when the time-scale gets larger (a week, say), the probability distributions start to resemble normal distributions with much lighter tails. Having a lighter tail means that there is a lower probability of an extreme event. Thus, eventually, the independence assumption breaks down and the ugliness (if a black swan should be regarded as such) is an artifact of the assumption.

We could approach the Flash Crash of May 6, 2010, in another way. By not assuming independence, we could let the observations at the highest-frequencies to be serially dependent. Then a negative price change would be more likely to be followed by another negative price change than a positive one. Using the terminology of Sornette (2009), events that grow extreme and surprising through such a dependent process are called “Dragon-Kings.”[22] According to Sornette (2009), events created this way can become more extreme than Black Swans. The advantage of this approach is that the sampling frequency becomes an important dimension in the analysis of extreme events while in the Black Swan approach it is assumed to be something quite ad hoc.

What this digression is meant to teach is that in the era of high-technology seemingly unrelated processes may become coupled at the very short time-scales, leading to surprising extreme events that could not have been easily predicted using knowledge from older technology. Conceptually, there is some amount of chaotic behavior involved in the sense that small changes in a variable can accumulate to unexpected end-results in a large scale. The importance of time (or time-scale) cannot be underestimated and it must be a central aspect of the analysis of an extreme event. This means that probability distributions should be modeled over various time-scales with special focus on the short time-scales if the activity of HFT is of prime interest.

Front running and predatory trading

The fact that volatility is positively correlated with HFT on any specific day, such as May 6, 2010, does not of course prove that this is generally so. If HFT resembles the activity of market makers, volatility may decrease, as discussed in the “Good” later. Nevertheless, some HFT strategies taking advantage of short-term correlations have earned an especially ugly reputation. Front-running strategies are often mentioned in media.[23] In a typical front-running strategy, a high-frequency trader steps in front of a large institutional order and drive the price to a direction that is beneficial for herself. This can be achieved by for example co-locating a computer server close to a stock exchange, making it possible to react very fast (in milliseconds, or less) to incoming

22 Technology Review. (2009). “How Dragon Kings Could Trump Black Swans.”
23 New York Times. (July, 2009). “Stock Traders Find Speed Pays, in Milliseconds.”

orders. Co-locating has become a profitable business to stock exchanges themselves, limiting their interests to monitor or regulate their clients more.[24] The situation is in some ways similar to what the NYSE specialists (nowadays called designated market makers) were able to do years ago. They were the ones who could see the state of the order book and could forecast the direction of short-term price movements with high probability. Thus, similar strategies have been used in the past. The difference is that now speed and competition have replaced the protected status of certain elite groups. The proponents of today's market structure would say that the game has gotten fairer.

Front-running is an example of predatory trading. Predatory trading, more generally, refers to strategically placed trades that hunt its "prey" by first trading in the same direction and then reversing the position and making a "kill." The most well known example of predatory trading, although not from the HFT world, is the fate of the Nobel Laureate lead hedge fund Long Term Capital Management (LTCM). In 1990s, LTCM was engaged in convergence arbitrage, a strategy in principle the same as the simpler pairs trading, in which the long-short positions were prepared to be taken for a long term (hence the name "Long Term"). Although quite sophisticated, the weak point of the strategy was that the LTCM positions were highly leveraged and hard to keep as a secret. In 1998, in the aftermath of the Russian debt crisis (default on its ruble nominated debt), the positions of LTCM in certain illiquid assets became well-known in many large Wall Street investment banks, such as Goldman Sachs. Many of these banks then started trading against the known LTCM positions and eventually drove LTCM to the brink of bankruptcy – which was, interestingly, a bit later re-capitalized by a consortium of some of the same banks that had been (claimed to be) the predators. Thus, predatory trading is nothing new to Wall Street.[25] The techniques to do it are now different, however, and rely much on technology, speed in particular. De Luca et al. (2011) report experimental evidence that it is speed, and not (artificial)

24 Wall Street Journal. (July, 2009). "NYSE's Fast-Trade Hub Rises Up in New Jersey."
25 Predatory trading is not new to finance or economics at large either. One of the culprits of the current debt crisis in Greece is said to be predatory traders that bet against Greece's survival in euro. [New Scientist (No. 2866, 2012).]

intelligence, that is the primary advantage of algorithms over humans – so far, that is.

“Predatory algorithms,” “immediate or cancel orders,” and “dark pool pinging” rely much on speed advantage. Essentially, they try to determine the state of large institutional orders by “sniffing them out” and using this knowledge to make (almost) risk-free arbitrage by trading on different exchanges in milliseconds or faster. Institutional investors using slower and not-so-smart algorithms lose. The industry standard is to execute large orders based on historical volume using the so-called volume weighted average price (VWAP) algorithm. Such algorithms are relatively easy to spot and predict by predatory HFT [see Agatonovic et al. (2012)], creating a technological adverse selection problem for institutional investors. This problem was, by some accounts, the key triggering mechanism of the Flash Crash of May 6, 2010.

Some of the best known predatory HFT strategies based on speed are called “quote stuffing”, “smoking”, and “spoofing.” Of these strategies, (quote) stuffing is arguably perhaps the most damaging to market quality. It limits the access of slower traders to markets by submitting a large number of orders and then canceling them very fast. This in turn leads to order congestion, which may in the worse case create technical trouble and quotes lagging by significant amounts of time. Egginton et al. (2012) find that over 74 percent of US listed equity securities experienced at least one stuffing episode in 2010 based on consolidated data from exchanges in the National Market System (NMS). They also find that during these episodes stocks experience decreased liquidity, higher trading costs, and increased short-term volatility. The smoking and spoofing strategies, on the other hand, try to manipulate other traders to participate in trading at an unfavorable moment, such as just before the arrival of relevant news.[26] Manipulative techniques are naturally not limited to these three strategies; for more examples of strategies, see for example BMO Capital Markets (2009). Most likely, smarter strategies are being developed, and the most ingenious ones kept as secrets.

26 For more details of manipulation techniques, see for example Biais and Woolley (2011) and Easley et al. (2012).

Predatory trading may damage liquidity provision from the side of slower traders. Because of the adverse selection problem, slower traders may not want to supply liquidity in the fear of being taken advantage of. However, evidence of predatory trading is still largely anecdotal. We know that high-frequency traders foresee the actions of slower traders to some extent, but it is not clear if this should be labeled predatory. Easley et al. (2012) state that it is obvious that the goal of many HFT strategies is to profit from slower traders' mistakes. Hirschey (2011) finds positive correlation between HFT and slower trading and this correlation is not simply due to faster reaction to arriving news by high-frequency traders. High-frequency traders seem to anticipate the buying and selling pressure of slower traders, increasing the trading costs of slower traders. Similarly, McInish and Upson (2012) find that fast traders are able to pick off slower liquidity demanders at prices that are not the best available. They estimate the profit to the faster traders to be more than $233 million annually and explain this trading opportunity to arise largely from how the US markets are constructed, and more specifically, due to the so-called "Flicker Quote Exception to Rule 611." But there is also evidence supportive of the contrarian view: Aldridge (2012) does not find any sign of HFT "pump-and-dump" gaming in futures markets. Thus, the verdict on HFT and their alleged predatory trading is not out yet.

Theoretical work on predatory trading is accumulating. In the model of Brunnermeier and Pedersen (2005), perhaps the first serious modeling framework applied to this subject, predatory trading leads to price overshooting and illiquidity when most needed. Furthermore, a liquidity crisis may spill over across traders and markets, affecting the financial markets as a whole and leading to a growth in systemic risk. Jarrow and Protter (2011) find that HFT can create mis-pricing disadvantageous to the slower traders. Decreasing profits to the slower traders is also forecasted by Hoffman (2011) who allows computers to trade with human traders. Most of the pot is taken by the faster (algorithmic) traders. The model of Biais et al. (2011) forecasts that increasing adverse selection costs lead to an arms race and excessive investment in HFT, because there is an increasing incentive to become a high-frequency trader. Pagnotta and Philippon (2012) find competing on speed to have negative welfare effects when the default speed reaches a threshold faster than is humanly possible. Brunnermeier and Pedersen (2005) suggest that batch auction markets, trading halts, and circuit breakers may be used alleviate the problem of price overshooting in their model. Policy recommendations are discussed at the end of the "Bad" section below.

We now move to the bit more abstract, but potentially worse arguments. The stand against HFT is reflected in the opinion of Nobel Laureate Paul Krugman: "The stock market is supposed to allocate capital to its most productive uses, for example by helping companies with good ideas raise money. But it's hard to see how traders who place their orders one-thirtieth of a second faster than anyone else do anything to improve that social function... we've become a society in which the big bucks go to bad actors, a society that lavishly rewards those who make us poorer."[27] The difference to the above discussion above is heavier emphasis on academic literature.

THE BAD

Academic literature on the "Bad" is surprisingly sparse, although it appears to be increasing. To the best of our knowledge, Zhang (2010) is the only credible study finding significant empirical evidence against HFT – that is, if event studies, such as Kirilenko et al. (2011), are excluded. The reason for the lack of evidence may be, as noted in the beginning, that not all studies make a distinction between HFT and AT (or other forms of automated trading) and use inaccurate proxies for them. This is understandable to some extent, because it is not easy to single out HFT from other trading and because such data have proprietary value. But it is somewhat surprising that Zhang (2010) finds evidence that contradicts the evidence of the "Good" section

27 New York Times. (2009). "Rewarding Bad Actors."

by only separating HFT from institutional and retail trading using common US databases. We thus review the results of this study in greater detail than otherwise.

Volatility

Zhang (2010) studies the long-term effect of HFT on volatility and if HFT aids or hinders the markets in reflecting the true prices, that is, if fundamental news are incorporated in market prices efficiently. According to the estimates presented by the author, HFT was responsible for about 78 percent of the dollar trading volume in 2009 in the US while in 1995 it was practically zero.[28] Zhang (2010) finds that HFT increase long-term stock market volatility. The effect is stronger for the largest 3000 market capitalization stocks and stocks with high institutional holdings. This finding is in agreement with the argument presented in the “Ugly” section that HFT firms take advantage of large trades of institutional investors. It is also in agreement with the theoretical results of Cartea and Penalva (2012) whose model predicts HFT to increase the price impact of large institutional investors making sizable portfolio changes. If their market impact is higher, it is natural to expect higher volatility too.

One could argue that the market impact effect would be most evident in short-term volatility. It does show in the model of Martinez and Roşu (2011). There is also empirical evidence of increased short-term volatility due to HFT. By analyzing an international dataset covering 39 exchanges worldwide, Boehmer et al. (2012) find that short-term volatility was systematically increased by AT/HFT, defined and measured following Hendershott et al. (2011) (see the “Good” section later). The authors argue that the increase in short-term volatility is hard to explain by faster absorption of news. They however do not explicitly consider the possibility that higher volumes (the effect of bid-ask bouncing, say) could be the main cause of it. Volatility created by microstructure would be flickering, reaction to news bursting.

28 Financial Times. (August, 2009). "High-Frequency Trading Under Scrutiny."

The empirical evidence concerning HFT and (short and long-term) volatility is controversial to some extent. But the fact that HFT leads to higher trading volumes cannot be disputed. Empirically, Dichev et al. (2011) establish a positive relationship between volume and volatility. If HFT leads to higher volumes and volumes lead to higher volatility, we have causality. Dichev et al. (2011) find that trading in general creates its own volatility beyond fundamental changes, consistent with larger market impacts and price overshooting. The theoretical model of Jarrow and Protter (2011) also forecasts higher volatility in the presence of high-frequency traders. However, one must stay critical of such modeling frameworks, because they make quite radical simplifying assumptions that may not hold in reality. For example, Jarrow and Protter (2011) assume frictionless and competitive markets where no bid-ask spread exist and the market is perfectly liquid.[29] Similarly, some empirical results of Zhang (2010) remain without sufficient proof. In particular, the found positive correlation between HFT and volatility when market uncertainty is high does not reveal the direction of causality. Volatility may be independent of HFT and be caused by something else. Causality is hard to establish, criticism that is repeated in the “Good” section later.

Price discovery and transparency

The second main finding of Zhang (2010) is that HFT is negatively associated with the market's ability to incorporate news about a firm's fundamentals into asset prices. Zhang (2010) finds that news about firm fundamentals lead to stock price overreaction when there is a lot of HFT volume in the markets. Because the reaction is almost totally reversed in the subsequent period suggests that HFT disturbs price discovery. Why this would be so can be justified by HFT firms' almost total neglect of fundamentally “true” (fair or intrinsic) value of companies. HFT firms are typically interested only in short-term dynamics, leaving valuation to analysts and

29 Another, more general, criticism is that these models typically categorize market participants quite carelessly as liquidity providers and takers, or informed and uninformed. This may be a too simplistic view of the real markets.

other “fundamentalists.” The overreaction hypothesis gets some theoretical support from Froot et al. (1992) who find that short-term speculators in general may put too much emphasis on short-term information and not enough on stock fundamental information, leading to a degradation of the informational quality of prices. It would not necessarily be a problem if the mis-pricing would disappear fast, but it becomes a problem if the mis-pricing catches momentum. In this sense, highly active short-term trading may turn out to be counterproductive from the view point of price discovery.

The claim that HFT is detrimental to price discovery may seem surprising at first as one might suspect that more trading volume would just lead to more precise pricing. But the increased trading volumes may be an illusion. Much of this trading volume is most likely out of reach of ordinary traders who are not fast enough to react to HFT. Grillet-Aubert (2010) reports that order cancellation by HFT firms can take place in less than 10 microseconds and that “message-to-trade” ratios tend to be extremely high for the most actively traded stocks: 96.5 percent of three HFT hedge funds' orders were reportedly cancelled (these three funds accounting for 39.6 percent of trading in the most liquid French stocks). Similar evidence is presented by Van Kervel (2012) who reports ratios of 31:1 and 51:1 for London Stock Exchange and and Chi-X, respectively. He finds that the liquidity offered of interlinked trading venues may not be realizable because some HFT market makers duplicate their limit order schedules on several different venues to increase their execution probabilities. When a trade is executed on one venue, the orders on others are quickly canceled.

Message-to-trades ratios are obviously related to market transparency. Does the visible state of the order book represent real tradable opportunities or not? If orders appear and disappear so fast that they cannot be reacted to by anyone else than computers – and more precisely, computers co-located near stock exchanges – it is not clear how to place your buy or sell order to get an exactly fair price. On the other hand, this only matters if you are solely interested in making a dollar trading ultrafast because otherwise you would simply accept a reasonable offer and keep the asset for the long-term. The bad thing here is, to repeat the point made above, that HFT can in a sense prevent perfect price discovery and only remove some arbitrage opportunities. Market transparency can also be argued to be hurt by HFT because most of the order submission and cancellation activity takes place at or very close to the best bid and ask prices, which are typically relatively small amounts. Thus, beyond the best levels, the order book can look rather sparse in relation to what it was in the "old days." One might suspect with good reason that in the era of ultrafast automatized trading it is more probable that a large order will need to "walk in the book" to get fully executed.

We must note, however, that the feature of large (or larger than average) trades being executed at prices worse than the best price is not really a HFT specific problem but related to decimalization, that is, the decrease of the minimum tick size or price change [see, e.g., Vuorenmaa (2010)]. In the US, the preliminary steps were taken in this direction already in the 1990s. First there was the move from the traditional Spanish "pieces of eight" ($1/8) to $1/16. Finally, in 2001, down to cents ($1/100). As a consequence, the sizes of orders decreased noticeably in the US stock markets. The completion of the decimalization process motivated large institutional investors to start using AT strategies, such as WVAP-based strategies, to minimize their impact on prices and to hide their trading intentions from the others to avoid being front-ran. Thus, retrospectively, the doomsday dynamics of the forthcoming flash crashes were set in motion by regulatory changes – RegNMS[30] in the US, possibly MiFID[31] in Europe – and the following decentralization of markets that fuel automated trading. There is no doubt that these regulatory changes have in part helped HFT to blossom.

More subtle changes

It is clear that regulatory and technological changes have changed the landscape of

30 Regulation National Market System (RegNMS).
31 Markets in Financial Instruments Directive (MiFID).

trading, and HFT activity has made the change even more dramatic. In addition to the aforementioned phenomena, there appear to be more subtle dynamic changes going on as well. This is not necessarily bad if the new price dynamics make the markets more efficient and healthier. We categorize the following subtle changes as “Bad” to be conservative and because they may affect the systemic risk as explained below.

First, the theoretical model of Cvitanic and Kirilenko (2010) demonstrates that in limit-order book markets populated by low-frequency human traders, the introduction of HFT machine traders leads to a distribution of transaction prices that has more mass around the center and thinner far tails. The authors argue that the shape of the transaction price density is consistent with HFT machines making positive expected profits by “sniping out” human orders somewhat away from the front of the book. The introduction of HFT also leads to shorter times between trades and higher volume, following from the assumption that HFT provides liquidity strategically.

Second, Smith (2010) uses wavelet methodology to estimate the Hurst parameter, which reflects the long-memory characteristics of a stochastic process. The author's main finding is that the Hurst parameter has increased substantially since the wider adoption of HFT after about 2005 and that the changes in the prices of some liquid stocks (and in their volume) have started to present fractal-like self-similarity at increasingly shorter time-scales – especially on the NYSE where automation is a newer phenomena compared to NASDAQ.[32] A change in the price dynamics at short time-scales is indeed plausible as HFT has become the dominant force at those time scales. While the study makes some doubtful arguments,[33] analyses a limited number of stocks, and does not separate HFT from AT, it is nevertheless interesting to find that the long-memory characteristics are increasing over time. In practice, this means that the traded volumes have become more predictable, which makes sense because

32 ZeroHedge. (2010). “Scientific Proof That High Frequency Trading Induces Adverse Changes In Market Microstructure And Dynamics, And Puts Market Fairness Under Question.”

33 It is for example incorrect to state that self-similarity has increased because pure random walk in continuous-time, that is, Brownian motion, is as self-similar as fractional Brownian motion with a higher ($H > 0.5$) Hurst exponent.

algorithmic traders, in particular, spread their trades over time. In a related study, Chaboud et al. (2011) show that AT/HFT strategies are more correlated over time than strategies operated by humans manually. Similarly, Brogaard (2010) finds that HFT strategies are more correlated with each other than non-HFT strategies. Such empirical results suggest that HFT may lead to increases in systemic market risk.

In addition to having more correlated asset prices especially at the intra-daily level, there is growing concern of correlation across assets and markets. This heightened correlation manifest itself partly in the Flash Crash of May 6, 2010, where different asset classes became quickly coupled through hedging attempts. The trend towards higher correlations in financial markets is likely to continue rather than to reverse. Finance Watch (2012) notes that cross-market arbitrage strategies performed on systematic, automated basis at high speeds may lead to worse contagion effects across markets. Fads in algorithms and inability to protect them intellectually can lead to quant-meltdowns described in Khandani and Lo (2009). Effectively, markets may become overcrowded and the population of algorithms too homogenous. Markets Committee (2011) note that the adoption of equity market trading technology to FX markets has created a more fertile ground for crashes to flourish there. This is worrying, because a crash in the FX markets could propagate deeper in the world economy. Debt and interest rate markets could also be dominated by algorithms. Haldane (2011) and Farmer and Skouros (2011) discuss stability issues in the era of automatized trading. We next review some of the proposed policy recommendations.

Proposed policy recommendations

The Bank of England's director of financial stability, Andrew Haldane, has become well known for his strong policy recommendations based on the Flash Crash of May 6, 2010 [see Haldane (2011)].[34] They are part of the “Foresight Project” of the

34 News release on a speech by Andrew Haldane, "The Race to Zero."

Government of England set to "explore how computer generated trading in financial markets might evolve in the next ten years or more" and its effect on financial markets more generally.[35] In Haldane's view, the most technologically-savvy traders have gained an unfair edge over the less sophisticated investors, which is comparable to being informed of the true value of a stock. The observed increasing correlations, and the potentially larger systemic risk lay at the heart of his worries. The "race to zero" in terms of speed increases systemic risk and has no winners. Haldane (2011) forecasts that the arms race of trading technologists will drive financial markets to the brink of a collapse unless stricter regulations are placed on HFT. This forecast is supported by the theoretical equilibrium prediction of Biais et al. (2011), mentioned in the "Ugly" section above, in which large institutions are fast and informed while small institutions are slow and incur adverse selection costs. "Grit in the wheels, like grit on the roads, could help forestall the next crash" is Haldane's policy proposal.

Haldane (2011) suggests more extensive data analysis from the part of regulators. This should actually be feasible because financial data are to be cleared through central counterparties and trade repositories. Larger datasets would allow regulators to identify the potential troublemakers. For example, the methodology proposed in Easley et al. (2011a,b) using order imbalances to predict how toxic the order flow is and what the optimal trading horizon is, could be used to prevent catastrophes. Their methodology has been shown to have predictive power over the Flash Crash of May 6, 2010. Haldane (2011) also suggests to stress-test financial markets thoroughly by simulation-based techniques as discussed by Cliff (2010). By changing the initial conditions in simulations and recording the results after each realization, an ensemble forecasting approach could be used to give estimated probabilities of extreme events.

In a related Foresight-article, Farmer and Skouras (2011) suggest an ecological perspective to developing real-time warning signals for systemic risk and building a

35 Foresight Project: The Future of Computer Trading in Financial Markets.

deeper understanding of markets based on large-scale simulations. Their claim is that predator-prey relationships exist in the markets, some of which are beneficial and some of which are not. More generally, agent-based simulations (ABS) could yield useful insight on how heterogeneity among market participants and small changes of parameters affect financial markets [see, e.g., Aloud et al. (2012a,b)]. ABS might also be combined with experimental economics, a field that studies how human agents operate in electronic markets [see Cliff (2011) and De Luca et al. (2011)]. After all, human-computer interactions are nowadays prevalent in many financial markets.

Finally, Haldane (2011) suggests that market microstructures in Europe could be modified by for example stricter market making guidelines and by the use of circuit-breakers. The most radical suggestion appearing in Haldane (2011), among others, is the use of resting rules for trades. This would presumably increase quoted bid-ask spreads and decrease liquidity in normal times, but hopefully save the markets from instabilities in the form of flash crashes. The industry does not greet this proposal with much enthusiasm for the obvious reasons. The proposal ignores the fact that most the problems related to HFT seem to be related to the US market microstructure.

The general idea that high liquidity should not be a top priority is nothing new in the field of economics. Finance Watch (2012) notes that John M. Keynes – one of the most influential economists of all time – himself thought that liquidity can divert the attention from the true value of an asset by creating an incentive for speculation. The high trading volumes experienced today could thus be interpreted as proof of more speculative behavior in today's financial markets. Finance Watch (2012) is quick to note, however, that higher trading volumes should not confused with higher liquidity. Volumes can be high but liquidity low, as the Flash Crash of May 6, 2010, has shown. The view of Finance Watch (2012) is highly critical of HFT and its policy suggestions even more extreme than that of Haldane (2011). They propose, for example, that HFT firms should provide the authorities an access to their trading algorithms on regular basis, a proposal that will no doubt raise much alarm (or disgust) in the HFT industry.

The underlaying problem above could be stated in the form of two questions: Is the HFT firms' trading behavior mostly speculative and if so, does it lead to more “Bad” than “Good” overall? It could be argued by many that speculation is acceptable if the “Good” win the “Bad.” Essentially, then, does HFT carry out a valuable function for example by market making, and this way make the markets to function better overall? The answer by the HFT industry to the latter question is positive and it is backed up by significant amount of empirical as well as theoretical evidence as we next discuss.

THE GOOD

The HFT industry is, quite naturally, biased towards speaking of the “Good” rather than the “Bad” or the “Ugly.”[36] For this reason their arguments must be interpreted with similar cautiousness as the critics speaking of the “Ugly” and “Bad” above. The following presents mostly academic empirical results concerning HFT and AT. There is currently much ongoing work in the field of HFT and AT, not least because of the politically and economically loaded attributes. We try to stay critical of the results.

Profitability and industry view

The view of the HFT industry is a good place to start. Essentially, HFT proprietary firms and hedge funds believe their trading style gives them better risk protection than more traditional trading styles because of the short intraday holding times of HFT. Market risk is further decreased by their extremely fast absorption of news content and by taking minimal overnight positions when important news could arrive.

With less risk, HFT firms claim to find better profit opportunities than traditional

36 See, for example, the recent speech given by Jim Simons, the founder of Renaissance Technologies: video at 48:20.

investors. Brogaard (2010) estimate that HFT generates trading profits of $2.8 billion annually. TABB Group (2009) estimates $7.2 billion in net profits in 2009.[37] Kearns et al. (2010) put an upper bound to $3.4 billion in the US equity markets in 2008, but believe that the true profits to be more modest. Some real-world examples, such as Renaissance Technologies, exemplify that extraordinary profits can be made.[38] A large chunk of these HFT firms' revenues are thought to be driven by speed of execution and low-latency capabilities to quote and cancel orders. Moallemi and Saglam (2011) show mathematically the importance of having low-latency in contemporaneous decision making situations. Many brokers and exchanges also require large volumes to give better fee structures, so to play the game profitably, one must have sufficiently low latency and high volume. However, Easley et al. (2012) note that there might be more to this than just speed: in their view, HFT is more about strategic order placing and operating in another clock, not a physical one like the ordinary traders, but a volume induced clock.[39] Interestingly, this presumed HFT's edge would not disappear by setting any speed limits on trading. The industry itself not only believes that HFT gives them good profit opportunities, but also that it gives them better risk protection and more efficient handling of multiple liquidity pools. In industry's view, their profit maximizing leads to more efficient markets through lower trading costs (especially quoted spreads), lower volatility, higher liquidity, transaction transparency, price discovery, and more diversified market interactions. This industry opinion is well captured in a prominent HFT firm's public commentary to the SEC.[40]

Trading volumes have increased dramatically over the last decade, as already stated. Credit Suisse (2010), for example, reports tripled daily volumes in the US equity markets where the share of HFT of daily volume is typically reported to be in the range 50 to 70 percentage, being a bit lower in Europe. Similar HFT participation rates have been documented for Australia as well [see Lepone and Mistry (2010)].

37 Wall Street Journal. (August, 2011). "A Wild Ride to Profits."
38 Renaissance Technologies' Medallion Fund: Performance Numbers Illustrated.
39 This idea goes again back to the 1960's: it was Mandelbrot and his authors who suggested to use “subordination.”
40 Tradeworx. (2010). “Public Commentary on SEC Market Structure Concept Release.”

These rates are also in line with recent theory [see, e.g., Martinez and Roşu (2011)]. But high volumes alone cannot be considered "Good" unless they serve some good purpose. The argument is that higher volumes imply faster price discovery and higher liquidity and thus better market quality. But as discussed in the "Bad" section above, higher volumes do not necessarily lead to these goals. To solve this apparent contradiction, we must first understand what liquidity is and how to measure it. The problem is that no universally accepted definition of liquidity exists. This is a serious hinderance in that it enables one to argue both for and against by defining liquidity conveniently. For the purposes of clarifying problems in liquidity measurement, we take a constructive approach and define liquidity to consist of three dimensions: spread, depth, and resiliency [as done in, e.g., Linton and O'Hara (2011)]. There is however no standard way to study liquidity and HFT empirically. Most empirical studies complement each other. We start the "Good" review considering the spread.

Liquidity

Quoted (bid-ask) spread is the difference between the bid and ask quotes in the book. There is overwhelming evidence that spreads have declined over the last decade or so. Spreads started to dramatically decline already in the end of 1990s due to decimalization, as noted in the "Bad" section. Trading algorithms have narrowed quoted spreads further [see, e.g., Castura et al. (2010), Credit Suisse (2010), and Haldane (2011)]. In particular, Menkveld (2012) finds a 50 percent decrease in quoted spreads on Chi-X Europe. Hasbrouck and Saar (2011) study "low-latency automated trading," that is, trading that reacts to market events in milliseconds. They find low-latency automated trading to be associated with lower quoted and effective spreads[41] on NASDAQ. Their study highlights that it is crucial to accurately define and identify HFT. Hasbrouck and Saar (2011) approximate HFT activity by "strategic runs" linked to submissions, cancellations, and executions – the tell signs of HFT. It

41 Effective spread is the signed difference multiplied by two between the trade price and the bid-ask midpoint prevailing at the time of order submission.

is yet to be verified how good this proxy is. Nevertheless, their results are consistent with all the other presently available empirical evidence concerning HFT and spreads.

The fact that spreads have narrowed seems to imply that the costs of trading have lowered as well. In addition to retail investors, institutional traders have most likely benefitted from these spread decreases as well – assuming they have kept up with the technological pace. In practice, this means use of transaction cost analysis (TCA) and AT in order to minimize adverse price impacts of large orders: slicing of large orders into small child orders and distributing them over different trading venues. OXERA (2011) indeed finds that prices and costs of using infrastructure providers have come down, as have the costs of using intermediaries. But in line with the above described use of TCA/AT, they also note that although the cost per transaction has fallen, the costs expressed in terms of value of trading have increased in some financial centers. Thus, it cannot be unanimously concluded that trading costs have decreased overall.[42]

Smaller spreads and higher automation have reduced the power of privileged groups, such as designated markets makers (see the "Ugly" section). Effectively, smaller spreads make it more difficult to earn money from the spread. The good news for the "honest" investors is that in the automated world there appears seems to be less room for manipulation in the form of front-running and other predatory strategies. Opinions to the contrary have been suggested as well. Zigrand et al. (2011) believe that the greater anonymity of computer based trading has made predatory trading easier. But this misses the fact predatory trading behavior is typically "penny-picking" requiring substantial technological investments and the field is getting more competitive each day. The bad news for HFT firms is, of course, that it is harder to gain competitive edge without enough technological savviness (see the "Bad" section). Overall, the documented spread decrease reflects both technological innovation and the more competitive playing field especially in market making, whether voluntary or not.

42 This depends crucially on how transaction costs are implemented, that is, if costs are transaction or volume based.

Theoretically, the observed spread decrease is understandable due to fast quoting and cancellation, which minimizes exposure to asymmetric information risk (adverse selection), a desirable feature in market making. High message-to-trade ratios mentioned earlier are instrumental in decreasing the spread to low levels. Fiercer competition should take care of excessively large market making or HFT returns. Interestingly, Martinez and Roşu (2011) show that as the number of high-frequency traders increase, they tend to make the markets more efficient and more stable. It is natural to expect that first-comers gain more than late-comers. In the end, one would not expect any particular participant to rise above others. Repeating the success of the early HFT adaptors, such as Renaissance Technologies, is going to be much harder now unless new significant innovations or regulations change the playing field again.

The two other dimensions of liquidity, depth and resiliency, need to be discussed too. Angel, Harris, and Spatt (2010), for example, present evidence that depth (usually defined as the sum of standing orders in the order book up to some reasonable level), has improved in the last decade. Although an increase in depth is intuitive, it is actually not clear-cut. Fast submission and cancellation of orders can take place at (or near) the best bid and ask, which would not contribute to depth beyond the high levels as described in the "Bad." But while the transparency of the order book may have diminished after decimalization and the rise of automated trading, most empirical studies still find the total effect on depth to be positive. Credit Suisse (2010), for example, reports large increases in quoted sizes since 2004. Hendershott and Riordan (2011) find that automated trading demands liquidity when spreads are narrow and depth is high, but not in the converse situation. HFT firms make a service to markets in supplying liquidity when needed and taking liquidity when not critical. The "Good" side of HFT is also reflected in the fact that one of the largest HFT firms today, GETCO Securities, became a designated market maker on the NYSE in 2010, giving some reassurance that HFT can be trusted to provide liquidity also in trouble.[43]

The third dimension of liquidity, resiliency, is the ability of the price to revert fast back to its original price after an abnormal shift in it, perhaps due to a large order. This dimension of liquidity is harder to measure accurately than the two others. It depends on the time-scale of investigation similarly to Black Swans: although the Flash Crash of May 6, 2010, is as an example of a loss of resiliency, it is in some sense only a temporary one – prices bounced back quite fast.[44] More generally speaking, the market impact of an order indicates how resilient the markets are to a sudden increase in trading. One common way to measure this effect is the so-called Amihud illiquidity measure. Using this measure at a daily frequency, Boehmer et al. (2012) find that AT/HFT is associated with decreases in Amihud illiquidity measure, which is consistent with the results for the other more traditional liquidity measures. Gsell (2008), on the other hand, uses a simulation-based technique to find, quite expectedly, that AT increases market impact when larger volumes are being traded. Research concerning resiliency is still rather scant and more should be conducted.

While not all of the dimensions of liquidity have been exhaustively examined in the literature, numerous studies report increases in overall liquidity due to AT/HFT. We refer the interested reader to a review of equity markets by Gomber et al. (2011) and of FX markets to Markets Committee (2011). We close the discussion on liquidity by briefly mentioning two often referenced academic studies using wider proxies for AT/HFT. First, using the NYSE quote dissemination change to “autoquoting” in 2003 as an exogenous instrument,[45] Hendershott et al. (2011) show that trading algorithms improve liquidity and enhance the informativeness of quotes. One potential weakness of their study is, however, that they cannot observe what trade is actually computer

43 Traders Magazine. (December, 2011). "Designated Market Making Alive and Well at NYSE."
44 This opinion, and how it compares to the 1987 market crash, is raised for example by Jim Simons: video at 50:56.
45 The change in market microstructure was a response to the decreased depth that happened due to decimalization. “Autoquoting” allowed algorithmic liquidity suppliers to quickly notice an abnormally wide quote and provide liquidity accordingly via a limit order. It did not affect the NYSE specialists, however, making it a good instrument.

algorithm generated and what is human generated, forcing them to use the rate of electronic message traffic as a proxy for AT/HFT. Similarly, Chaboud et al. (2011) resort to a proxy showing the amount of human and computer generated trades (at one-minute frequency) and quotes (at one-second frequency) in the FX markets. Chaboud et al. (2011) do not find evidence of AT/HFT causing any liquidity shortages in the three currency pairs they analyze. But as both empirical studies note, the periods in their analysis are not truly tumultuous, making it uncertain how automated trading would behave in a true crisis period instead of a normal period.

Price efficiency and price discovery

Liquidity is closely related to the concepts of price efficiency and price discovery. There exist several definitions of price efficiency, but the most applied states that there cannot be any predictability in prices to be taken advantage of because they already reflect all relevant information. This, in turn, means that the true value of an asset is discovered correctly and fast. Increased HFT market making activity should in principle increase price efficiency and discovery. Increased predatory trading and manipulation could however swamp these benefits, but as noted for example by Linton and O'Hara (2011), we can expect competition to eliminate most of the adverse effects. More colorfully put, as the easy prey (slower traders) becomes harder to catch, predators turn to hunting other predators, which takes care of "excess rents."

Most of the empirical evidence in the literature find that HFT has an overall beneficial effect on price efficiency [see, e.g., Castura et al. (2010), Brogaard (2010), and Hendershott (2011)]. More specifically, Hendershott and Riordan (2012) conclude that HFT increase price efficiency by trading in the direction of permanent price changes and in the opposite direction of transitory pricing errors. The efficiency enhancing activities of high-frequency traders play a greater role in trading than any alleged manipulative strategies. In a related study, Hendershott and Riordan (2011) find that algorithmic traders' quotes are actually more efficient than other traders' quotes and their demand of liquidity moves the observed prices towards the efficient price. The limitation of the latter study (compared to the former) is however that it does not separate AT from HFT – a persistent problem of many empirical studies.

Volatility

That HFT increases liquidity suggests that volatility is dampened in the process. There is a good amount of empirical evidence showing that volatility and HFT are inversely related to each other: HFT tends to be low with high volatility and high with low volatility. Credit Suisse (2010), for example, reports dampening of short-term volatility "that might otherwise be created by large institutional orders filled during the day." In the academic side, Brogaard (2012) finds HFT to be associated with lower intra-day volatility on NASDAQ and BATS. Chaboud et al. (2011) do not find evidence of volatility increases in the FX markets. However, as described above, they use sporadically sampled data where increases in microstructure noise may not be so evident, as noted by Cartea and Penalva (2012). Hasbrouck and Saar (2011) find declined volatility in association to low-latency automated trading in equities.

An important question is do the HFT firms participate only in low-volatility regimes or do they actually lower volatility by their trading practices? Causality remains still largely unclear. Brogaard (2012) addresses the question econometrically using Granger-causality tests and a dataset that identifies over twenty HFT firms by their trading styles but cannot determine the direction of causality. Brogaard (2012) finds, however, that reducing HFT activity by the Short Sale Ban (in effect in September, 2008) increased intra-day volatility, suggesting that if regulations are placed to limit HFT, it could have adverse consequences in the quality of markets. In a related study, Gsell (2008) also finds that low-latency in AT decreases volatility using simulations. This gives some support to the controversial co-location practice of many exchanges.

The empirical findings concerning “Good” can be criticized by the fact that most of them are found in normal data conditions. This point is often repeated by regulators. In a speech given in May, 2010, the Chairman of SEC, Mary L. Schapiro, questions whether HFT algorithms are "programmed to operate properly in stressed market conditions."[46] Observations such as the Flash Crash of May 6, 2010, are statistically speaking infrequent events and no scientifically convincing evidence is yet available. There is at least some reason for concern, as shown recently by Boehmer et al. (2012) using an extensive international database. By identifying difficult market making days to experience exceptionally large two-day cumulative returns, they find that some of the “Good” aspects are decreased and short-term volatility is increased more by AT/HFT than in normal times. They however again have to resort to using a proxy following Hendershott et al. (2011). In any case, the results suggests that HFT firms avoid market making on days that put them in disadvantage in relation to informed investors. This is understandable as most of the HFT firms do not have a designated market making role and the markets are highly competitive. Moreover, other positive effects of HFT persist on difficult days, such as liquidity, although in reduced amount.

It is possible to increase the statistical significance by analyzing more frequent events and hope that this would shed some more light on extreme events. Brogaard (2012) does this by studying the effect of macro and company specific news on HFT activity and volatility. By combining different datasets, Brogaard (2012) finds that HFT firms tend to take liquidity after macroeconomic news announcements and supply liquidity after company specific news. Macroeconomic news announcements are scheduled and much less frequent than randomly arriving company specific news. Most of the empirical evidence of Brogaard (2012) supports the view that HFT reduces volatility. This would also seem to be consistent with the theoretical predictions of Jovanovic and Menkveld (2011) who find that HFT can increase welfare by quickly updating

46 Schapiro, Mary L. (2010). "Strengthening Our Equity Market Structure."

quotes on news in general (they do not consider different type of news separately) and thereby decrease adverse selection costs on price quotes. However, Jovanovic and Menkveld (2011) do not consider abnormal time periods in their analysis either.

Proposed policy recommendations

While the SEC is currently contemplating on how to regulate HFT, it is interesting to note that its former chairman, Arthur Levitt Jr., takes a strong stand in favor of HFT.[47] His argument is based on the competitiveness and health of the US markets. HFT is, in his view, only the most recent stage in technological innovation and authorities should thus not "regulate out of existence all high-frequency trading." There is much truth to this claim in that restricting progress has never worked out well in human history. The insights offered by Easley et al. (2012) support this claim: HFT is not so much about speed than believed, but more about strategic order placing and being smarter than others. Instead of suggesting stricter speed limits for HFT, they offer solutions for the slower traders not to be easy targets for potentially predatory HFT algorithms. Similarly, one could try to create more detailed datasets for regulators by following the suggestion of Donefer (2010): in particular, mark all automated trades as belonging to one of the four types mentioned at the beginning of this article: (i) liquidity seekers; (ii) market making; (iii) statistical arbitrage, or (iv) rebate seeking.

Technologically, as mentioned earlier, it does not make sense to put the brakes on because the same technological advancement could later on prove beneficial in other contexts. Stricter regulations or fee structures on HFT could also be suboptimal. They could lead to substantial welfare losses by shifting HFT from being mainly a liquidity provider to an aggressive liquidity consumer, as predicted by a model of Jovanovic and Menkveld (2011). The problem that SEC and its cousin regulatory authorities in Europe face is that they need to define HFT in order to regulate it. As of yet, they

47 Wall Street Journal. (August, 2009). "Don't Set Speed Limits on Trading."

have not accomplished to do this. With respect to the quality of the markets, Levitt argues that there is nothing wrong in looking for inefficiencies and exploiting them fast, especially if this facilitates liquidity, pricing, and healthy transparent markets.

An introductory book on HFT by Aldridge (2011) summarizes a few key social benefits of HFT: increased market efficiency, added liquidity, innovation in computer technology, and stabilization of market systems. These, and several other, benefits are supported by academic studies as reviewed above, although there is some concern on how stable the markets are with HFT. Although for example Hendershott and Riordan (2012) find no evidence of stability problems, the fear is that HFT market makers would disappear when they are most urgently needed. One solution to this could be to require market makers to always trade. Venkataraman and Waisburd (2007) show that there are benefits of having designated (but not otherwise privileged as earlier on the NYSE) market makers directly compensated by the listing firm. However, history has taught us that even designated market makers prefer to stop providing liquidity in difficult times: they rather face penalty fees than go broke.[48] As a potentially more flexible solution, real-time variable “make-and-take” fees have been suggested to attract liquidity. Effectively, in times of high market stress, larger rebates would be paid out to liquidity providers to attract them. The theoretical model of Foucault et al. (2012) could form a basis for a pre-analysis. Rebate systems of some sort are already in existence in many exchanges, so they would not be a very difficult thing to change.

In addition to the stability question, the now-heated battle over the “Good” and the “Bad/Ugly” aspects of HFT is driven by the confusion on defining central concepts. To solve this situation, we need to pay more attention to what we are talking about. First, we need a good working definition of HFT and reach a consensus also of how to measure liquidity and other key concepts for HFT. Several characteristics of HFT have been suggested in the literature, but no all-satisfying definition is yet found.

48 Katz, Gary. (2011). "Relying on the 'Rules of the Game' in the U.S. Options Market."

Pressure to find it is building up among regulators: CFTC Commissioner Scott D. O'Malia, has suggested a seven part test for what constitutes HFT.[49] Second, we need extensive data of HFT covering different time periods for better statistical analysis. We need to identify HFT and separate it from other forms of electronic trading, an issue often emphasized in the literature as well [see, e.g., Boehmer et al. (2012)]. Until then, we are speculators placing probabilities on “Good,” “Bad,” and “Ugly.” We hope the reader is now in a more favorable position to place a bet on her own.

CONCLUSIONS

We have reviewed the current literature on “the Good, the Bad, and the Ugly” aspects of HFT. Based on the results of mostly academic studies, both of theoretical and empirical nature, we are unfortunately not able to make a definite conclusion which interpretation is the most correct one from a neutral standpoint. It seems, however, that the “Good” aspects have the most weight and dominate the other aspects. This said, evidence on the “Bad” and the “Ugly” are gaining momentum in the literature. Once we have clearer working definitions for HFT and AT and gain better access to more exhaustive datasets we should be able to make sound policy recommendations.

To decompose this general conclusion in a more detail and usefully, we state the main message of each category in turn. First, the main message of the “Ugly” is that speed kills: flash crashes and other extreme events eventually force financial markets down on their knees by higher systemic risk that is HFT induced and, in the process, wealth is being redistributed from ordinary traders to predatory-like high-frequency traders. Second, the main message of the “Bad” is that HFT is a Pandora's Box: the seemingly good attributes of HFT are a hoax and market quality in the form of volatility, liquidity, transparency, and price discovery is mutilated and the health of the markets jeopardized. Finally, the main message of the “Good” is that progress should not be

49 O'Malia, Scott D. (2011). "Letter to the Technology Advisory Committee regarding the definition of High Frequency Trading."

denied: HFT is part of the 21th century trading technology and acting competitively it guarantees better performing financial markets in about every respect imaginable.

To be fair, each of these arguments are excessively pointed. The truth lies somewhere in the midst of the three dimensions. Obviously, the “Bad” and the “Ugly” are closer to each other in the three dimensional space than the orthogonal “Good.” The lesson that might be learned from our literature review and discussion is that policy and regulations will play an increasingly large role in the future of financial markets. Looser regulations have allowed the fragmented, competitive structure we see today. Now they have the capacity to reform the financial markets again. HFT, and possibly other forms of automated trading, should present their case clearly and loudly, so that when the regulations eventually get changed, the health of the markets is actually improved and not impaired. Technological advancements in the form of better simulation capability of extreme scenarios could be used as a general guideline here. Empirical experiences from the US and Europe should be evaluated rigorously and their differences and flaws compared in a commensurate way. There is no turning back to age-old trading practices: technology is here to stay and we must control it.

REFERENCES

Agatonovic, Milos, Vimal Patel, and Chris Sparrow. (2012). "Adverse Selection in a High-Frequency Trading Environment." *Journal of Trading* 7, 18–33.

Aldridge, Irene. (2012). "Can High-Frequency Traders Game Futures?" *Journal of Trading* 7, 75–82.

Aldridge, Irene. (2010). *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems.* Wiley Trading.

Aloud, Monira, Maria Fasli, Edward Tsang, Alexandre Dupuis, and Richard Olsen. (2012b). "Modelling the High-Frequency FX Market II: a Systematic Exploration of Market Features." Working Paper, University of Essex.

Aloud, Monira, Maria Fasli, Edward Tsang, Alexandre Dupuis, and Richard Olsen. (2012a). "Modelling the High-Frequency FX Market I: an Agent-Based Approach." Working Paper, University of Essex.

Angel, James J., Lawrence E. Harris, and Chester S. Spatt. (2010). "Equity Trading in the 21st Century." Working Paper, USC Marshall School of Business.

Bertsimas, Dimitris, and Andrew W. Lo. (1998). "Optimal Control of Execution Costs." *Journal of Financial Markets* 1, 1–50.

Biais, Bruno, and Paul Woolley. (2011). "High-Frequency Trading." Working Paper, Toulouse School of Economics and London School of Economics.

Biais, Bruno, Thierry Foucault, and Sophie Moinas. (2011). "Equilibrium High Frequency Trading." Working Paper, Toulouse School of Economics and HEC.

BMO Capital Markets. (2009). "The Impact of High Frequency Trading on the Canadian Market." Quantitative Execution Services Report.

Boehmer, Ekkehart, Kingsley Y. L. Fong, and Julie Wu. (2012). "International Evidence on Algorithmic Trading." Working Paper, EDHEC Business School, Australian School of Business, and University of Georgia.

Brunnermeier, Markus K., and Lasse Pedersen. (2005). "Predatory Trading." *Journal of Finance* 60, 1825–1863.

Brogaard, Jonathan A. (2010). "High Frequency Trading and Its Impact on Market Quality." Working Paper, Northwestern University.

Brogaard, Jonathan A. (2012). "High Frequency Trading and Volatility." Working Paper, University of Washington.

Carlson, Mark. (2006). "A Brief History of the 1987 Stock Market Crash – with a Discussion of the Federal Reserve Response." Working Paper, Board of Governors of the Federal Reserve.

Cartea, Álvaro, and Josè Penalva. (2012). "Where is the Value in High Frequency Trading?" Working Paper, Universidad Carlos III de Madrid.

Carlin, Bruce Ian, Miguel Sousa Lobo, and S. Viswanathan. (2005). "Episodic Liquidity Crises: Cooperative and Predatory Trading." *Journal of Finance* 42, 2235–2274.

Castura, Jeff, Robert Litzenberger, Richard Gorelick, and Yogesh Dwivedi. (2010). "Market Efficiency and Microstructure Evolution in U.S. Equity Markets: A High-Frequency Perspective." RGM Advisors, LLC.

CFTC/SEC. (2010). "Findings Regarding the Market Events of May 6, 2010." Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.

Chaboud, Alain, Erik Hjalmarsson, Clara Vega, and Ben Chiquoine. (2011). "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market." Discussion Paper, FRB International.

Chakravarty, Sugato, Pankaj K. Jain, Robert Wood, and James Upson. (2009). "Clean Sweep: Informed Trading Through Intermarket Sweep Orders." Working Paper, Purdue University, University of Memphis, and University of Texas at El Paso.

Cliff, Dave. (2010). "The Flash Crash of May 6th 2010: Why Technology Failures?" Speech, Royal Society of London.

Cliff, Dave. (2011). "The Impact of Technology Developments." Working Paper, Foresight.

Cliff, Dave, Dan Brown, and Philip Treleaven. (2011). "Technology Trends in the Financial Markets: A 2020 Vision." Working Paper, Foresight.

Connell, Liam, Richard B. Gorelick, Adam Nunes, and Cameron Smith. (2010). "Re: Concept Release on Equity Market Structure (File No. S7-02-10)." Letter to the SEC.

Credit Suisse. (2010). “Sizing Up US Equity Microstructure.” AES Analysis, Market Commentary.

Cvitanic, Jaksa, and Andrei Kirilenko. (2010). “High Frequency Traders and Asset Prices.” Working Paper, California Institute of Technology.

De Luca, Marco, Charlotte Szostek, John Cartlidge, and Dave Cliff. (2011). “Studies of Interactions Between Human Traders and Algorithmic Trading Systems.” Working Paper, Foresight.

Dichev, Ilia D., Kelly Huang, and Dexin Zhou. (2011). “The Dark Side of Trading.” Working Paper, Emory University School of Law.

De Jong, Frank, and Barbara Rindi. (2009). *The Microstructure of Financial Markets.* Cambridge University Press.

Donefer, Bernard S. (2010). “Algos Gone Wild: Risk in the World of Automated Trading Strategies.” *Journal of Trading* 5, 31–34.

Easley David, Lopez de Prado Marcos M., and O’Hara, Maureen. (2012) “The Volume Clock: Insights into the High Frequency Paradigm.” *Journal of Portfolio Management*, forthcoming.

Easley, David, Lopez de Prado Marcos M., and O’Hara, Maureen. (2011a). “The Microstructure of the “Flash Crash”: Flow Toxicity, Liquidity Crashes, and the Probability of Informed Trading.” *Journal of Portfolio Management* 37, 118–128.

Easley, David, Lopez de Prado Marcos M., and O’Hara, Maureen. (2011b). “The Exchange of Flow Toxicity.” *Journal of Trading* 6, 8–13.

Egginton, Jared F., Bonnie F. Van Ness, and Robert A. Van Ness. (2012). “Quote Stuffing.” Working Paper, University of Mississippi.

Fama, Eugene F. (1965). “The Behavior of Stock-Market Prices.” *Journal of Business* 38, 34–105.

Farmer, J. Doyne, and Spyros Skouras. (2011). “An Ecological Perspective on the Future of Computer Trading.” Working Paper, Foresight.

Farmer, J. Doyne, and John Geanakoplos. (2008). “The Virtues and Vices of Equilibrium and the Future of Financial Economics.” Discussion Paper, Cowles Foundation.

Finance Watch. (2012). "Investing Not Betting: Making Financial Markets Serve Society." Position Paper on MiFID/MiFIR.

Foucault, Thierry, Ohad Kadan, and Eugene Kandel. (2012). "Liquidity Cycles and Make/Take Fees in Electronic Markets." *Journal of Finance*, forthcoming.

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein. (1992). "Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation." *Journal of Finance* 47, 1461–1484.

Gomber, Peter, Björn Arndt, Marco Lutat, Tim Uhle. (2011). "High-Frequency Trading." Research Report, Goethe Universität.

Grillet-Aubert, Laurent. (2010). "Equity Trading: A Review of the Economic Literature for the Use of Market Regulators." Working Paper, Autorité des Marchés Financiers.

Gsell, Markus. (2008). "Assessing the Impact of Algorithmic Trading on Markets: A Simulation Approach." Working Paper, Center for Financial Studies.

Haldane, Andrew G. (2011). "The Race to Zero." Speech, Bank of England.

Hasbrouck, Joel, and Gideon Saar. (2011). "Low-Latency Trading." Working Paper, Johnson Graduate School of Management.

Hendershott, Terrence. (2011). "High Frequency Trading and Price Efficiency." Working Paper, Foresight.

Hendershott, Terrence, and Ryan Riordan. (2012). "High Frequency Trading and Price Discovery." Working Paper, Haas School of Business.

Hendershott, Terrence, and Ryan Riordan. (2011). "Algorithmic Trading and Information." Working Paper, NET Institute.

Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. (2011). "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66, 1–33.

Hirschey, Nicholas H. (2011). "Do High-Frequency Traders Anticipate Buying and Selling Pressure?" Working Paper, University of Texas at Austin.

Hoffman, Peter. (2011). "A Dynamic Limit Order Market with Fast and Slow Traders." Working Paper, European Central Bank.

Kearns, Michael, Alex Kulesza, and Yuriy Nevmyvaka. (2010). "Empirical Limitations on High Frequency Trading Profitability." *Journal of Trading* 5, 50–62.

Khandani, Amir E., and Andrew W. Lo. (2011). "What Happened to the Quants in August 2007? Evidence from Factors and Transactions Data." *Journal of Financial Markets* 14, 1–46.

Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. (2011). "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." Working Paper.

Jarrow, Robert, and Philip Protter. (2011). "A Dysfunctional Role of High Frequency Trading in Electronic Markets." Working Paper, Cornell University and Columbia University.

Johnson, Neil, Guannan Zhao, Eric Hunsader, Jing Meng, Amith Ravindar, Spencer Carran, and Brian Tivnan. (2012). "Financial Black Swans Driven by Ultrafast Machine Ecology." arXiv.org

Jovanovic, Boyan, and Albert J. Menkveld. (2011). "Middlemen in Limit-Order Markets." Working Paper, New York University and VU University of Amsterdam.

Lepone, Andrew, and Mitesh Mistry. (2010). "The New Breed of Market Participants: High-Frequency Trading Evidence from the Australian Stock Exchange." Working Paper, The Futures Industry Research Centre.

Linton, Oliver, and Maureen O'Hara. (2011). "The Impact of Computer Trading on Liquidity, Price Efficiency/Discovery and Transaction Costs." Working Paper, Foresight.

MacKenzie, Donald. (2011). "How to Make Money in Microseconds." London Review of Books.

Mandelbrot, Benoît B. (1963). "The Variation of Certain Speculative Prices." *Journal of Business* 36, 394–419.

Markets Committee. (2011). "High-Frequency Trading in the Foreign Exchange Market." Research Report, Bank for International Settlements.

Martinez, Victor H., and Ioanid Roşu. (2011). "High Frequency Traders, News and Volatility." Working Paper, Baruch College and HEC Paris.

McGowan, Michael J. (2010). "The Rise of Computerized High Frequency Trading: Use and Controversy." Duke Law and Technology Review.

McInish, Thomas, and James Upson. (2012). "Strategic Liquidity Supply in a Market with Fast and Slow Traders." Working Paper, University of Memphis and University of Texas at El Paso.

Menkveld, Albert J. (2012). "High Frequency Trading and the New-Market Makers." Working Paper, VU University of Amsterdam.

Merton, Robert K. (1949). *Social Theory and Social Structure.* Free Press.

Moallemi, Ciamec C., and Mehmet Saglam. (2011). "The Cost of Latency." Working Paper, Columbia University.

Morris, Stephen, and Huyn Song Shin. (2004). "Liquidity Black Holes." *Review of Finance* 8, 1–18.

O'Hara, Maureen, and Mao Ye. (2011). "Is Market Fragmentation Harming Market Quality?" *Journal of Financial Economics* 100, 459–474.

OXERA. (2011). "Monitoring Prices, Costs, and Volumes of Trading and Post-Trading Services." Report Prepared for European Commission DG Internal Market and Services.

Pagnotta, Emiliano, and Thomas Philippon. (2012). "Competing on Speed." Working Paper, New York University.

Perron, Charles. (1984). *Normal Accidents: Living with High-Risk Technologies.* Basic Books.

Securities and Exchange Commission. (2010). "Concept Release on Equity Market Structure." Release Number 34-61358.

Smith, Reginald. (2010). "Is High-Frequency Trading Inducing Changes in Market Microstructure and Dynamics?" Working Paper, Bouchet Franklin Institute. arXiv.org

Sornette, Didier. (2009). "Dragon-Kings, Black Swans, and the Prediction of Crises." arXiv.org

TABB Group. (2009). "US Equity High Frequency Trading: Strategies, Sizing, and Market Structure." Research Report.

Taleb, Nassim Nicholas. (2007). *The Black Swan*. Random House.

Van Kervel, Vincent. (2012). "Liquidity: What You See Is What You Get?" Working Paper, Tilburg University.

Venkataraman, Kumar, and Andrew C. Waisburd. (2007). "The Value of the Designated Market Maker." *Journal of Financial and Quantitative Analysis* 42, 735–758.

Vuorenmaa, Tommi. (2012a). "A Short Introduction to Automated High-Frequency Trading." Working Paper, Valo Research and Trading.

Vuorenmaa, Tommi. (2012b). "Liquidity, Activity, and Dependence on Interlinked Trading Venues." Revised Working Paper, Valo Research and Trading.

Vuorenmaa, Tommi. (2010). "Decimalization, Realized Volatility, and Market Microstructure Noise." Revised Working Paper, Valo Research and Trading.

Wood, Robert, James Upson, and Thomas H. McInish. (2012). "The Flash Crash: Trading Aggressiveness, Liquidity Supply, and the Impact of Intermarket Sweep Orders." Working Paper, University of Memphis and University of Texas at El Paso.

Zhang, Frank. (2010). "High-Frequency Trading, Stock Volatility, and Price Discovery." Working Paper, Yale University.

Zhang, Frank, and Stuart Baden Powell. (2011). "The Impact of High-Frequency Trading on Markets." *CFA Magazine* 22.

Zigrand, Jean-Pierre, Dave Cliff, Terrence Hendershott. (2011). "Financial Stability and Computer Based Trading." Working Paper, Foresight.