Government Office for Science

Foresight

Review of the benefits of a continuous market vs. randomised stop auctions and of alternative Priority Rules (policy options 7 and 12)

Economic Impact Assessment EIA11

Foresight, Government Office for Science

Contents

Review of the benefits of a continuous market vs. randomised stop auctions and of alternative Priority Rules (policy options 7 and 12)[1]

European Commission Public Consultation

Review of the Markets in Financial Instruments Directive

J. Doyne Farmer

Santa Fe Institute

jdf@santafe.edu

Spyros Skouras

Athens University of Economics and Business

skouras@aueb.gr

28 March 2012

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation.

[1] We note there is substantial overlap between Sections 3 and 4 of this Regulatory Impact Appraisal and another we have written (Farmer and Skouras 2011b) since our framework for thinking about both types of policies is related to the role of market speed. This Regulatory Impact Appraisal is more comprehensive in our treatment of those issues where there is overlap.

1. Objective

Financial market observers and stakeholders have expressed concern with the rapid acceleration in the speed of trading during the last decade. The “flash crash” of 2010 and several prominent technological failures of electronic markets make it important to question whether current trends in market design are safe and whether there exist better alternatives. We discuss a measure that could replace the currently entrenched market design with an alternative that might reduce market instability, reduce incentives for the relocation of financial market centers, increase competition, facilitate transparency and oversight and channel society’s resources to more productive activities.

It is important to emphasize that this is not a case of replacing a “market outcome” with a regulatory intervention: the currently entrenched market model has been strongly shaped by previous regulatory policies and historical accident in addition to market forces. Avoiding new regulations in this area simply means letting outdated regulations shape market outcomes.

Replacing the continuous double auction market model with a frequent pro rata sealed bid randomized auction model could have benefits that we would value in the billions of dollars.

2. Background

In order to understand the relative merits of continuous trading and randomized stop auctions and the relevance of previous studies in this area, it is necessary to first appreciate current market design and the role speed plays in it.

Currently the most widely-used trading mechanism in financial markets is the “continuous double auction electronic order book with time priority”. With this mechanism, quote arrival and transactions are continuous in time and execution priority is assigned based on the price of quotes and their arrival order (in what follows we refer to this as the “electronic order book” for short, even though this terminology normally encompasses other designs as well). In what follows, we focus our discussion on electronic order books for liquid instruments, which dominate stock and futures markets, but also appear in certain FX, options and bond markets.

One consequence of this trading mechanism which we believe is underappreciated is that it endows a huge advantage to being faster than other traders, creating evolutionary pressures that drive an arms race for ever-more speed. In a separate regulatory impact appraisal (Farmer and Skouras, 2011) we argue in detail why we believe market speed hinders the key job of markets to discover fair and stable prices and therefore why society might be significantly better off if it is reduced. We realize that these statements are controversial and that given the available data and the current state of the literature they may seem strong so we discuss them in great detail in Section 3. There we also sketch a model from which we can estimate the private value of relative speed to individual traders, i.e. from being just a little faster than a competitor.

The wedge between the impact of speed on society versus individuals is quite natural. Individuals benefit from being faster than other individuals which however has no benefit per se to society and indeed is probably harmful because it leads to a wasteful arms race and winner’s curse dynamics. This situation can easily be changed: The quest for speed is a side effect of trading protocol and can be altered by changing the protocol. Our appraisal will focus on a comparison of the currently used electronic order book against an alternative design featuring frequent pro-rata sealed bid call randomized auctions. The alternative design we propose is tailored to drastically reduce the private value of relative speed while at the same time being minimally disruptive to the functioning of price discovery and liquidity provision in markets.

If there were sufficient regulatory will do to so, our proposal - or at least its most important features - could be implemented and tried. Indeed, the work of Economides and Schwert (1995) is widely considered as having been influential in the adoption of opening and closing call auctions in most continuous trading markets. It is worth emphasizing that their arguments for call auctions revolved around the role of infrequent auctions in focusing markets' attention to price formation at key points in time, whereas we suggest frequent sealed bid auctions as a mechanism to slow down markets relative to current speed levels.

In the last few decades, speed and its incarnation as high frequency trading (HFT) has dominated the markets and the need to make trades quickly has become accepted as essential in financial markets. According to the Securities and Exchange Commission, HFT "is a dominant component of the current market structure and is likely to affect nearly all aspects of its performance". A variety of other sources estimate that HFT accounts for 60-70% of trading in the US and 30-40% in Europe.[2] Strikingly, only a tiny fraction of trading firms are active in high frequency trading. Iati (2009), for example, estimates that only 2% of the 20,000 trading firms operating in the US engage in HFT. As we will argue here, given current trading protocols it is not surprising that HFT is so dominant – there is a lot of money to be made. It is also not surprising that it is heavily concentrated in only a few firms – speed requires elaborate and expensive infrastructure, which must be continually improved to stay ahead of the competition. Expertise in such technologies seems orthogonal to the ability to make good investments and not all investors can or want to build and supervise HFT operations. The dominant trading protocol is shifting the skill-set of market participants away from that associated with long-term evaluations of the value of tradable securities.

2.1. The speed of today's markets

Information in today's networks travels at up to 75% of light-speed.[3] It takes approximately 70 milliseconds (hereafter mSec) for light to travel across antipodal points on the globe. New York to London is 18mSec, New York to Chicago about 4mSec. The fastest data connections between New York and Chicago currently have a latency of around 6.5mSec because of a combination of network imperfections and the fact that networks do not follow a straight line, but they are approaching their theoretical limit.

Markets also compete in terms of the speed of their matching engines, i.e. the time it takes to process an incoming order, calculate the state of the market and report it back to clients. The cutting edge is 124 microseconds (or millionths of a second) for the London Stock Exchange's

[2] The SEC quote is found in the SEC Concept Release, January 21, 2010. The TABB group reports HFT accounts for 61% of US equity trading volume (Tabb, Iati and Sussman, 2009) while other estimates exceed 70%.[2] Arnuk and Saluzzi (2009) report that HFT accounts for 60% of all volume on US equity, Zhang and Powell (2011) report that HFT accounts for 70% of consolidated volume in the US, 77% in the UK and 40% across Europe. AFM (2011) also reports estimates around 30-40% for European markets.

[3] See NYSE Euronext's report at www.latencystats.com/whitepaper.

Millennium matching engine which was introduced in 2011. Other exchanges such as BATS, NASDAQ, Chi-X, NYSE and TSE Arrowhead are all below 5 mSec.[4]

On the trader side, message processing is down to average latencies in the single digit microseconds with single digit standard deviations or "jitter".[5] As noted in AMF (2011), the life cycles of orders in European equity markets can be as short as seven microseconds. AMF (2011) reports that "low latency" strategies have latencies of 40 microseconds, "market making" strategies have 180 microseconds and "statistical arbitrage" has 200 microseconds. Hasbrouck and Saar (2011) find that 37% of all orders are cancelled within two seconds of being submitted in US equity markets.

It is widely agreed that nanosecond (billionths of a second) processing will soon be a reality and indeed, the technology exists from military and other industrial applications and will no doubt readily be adopted as soon as the business justification for doing so arises.

Some commentators observe that the speed of light poses a natural barrier on any trend in HFT and conclude that since this barrier is fast approaching, trends in HFT cannot continue for very long or at least not indefinitely. This is certainly true if one looks at absolute trends. However, as common sense suggests and as we will argue in the next section, HFT benefits almost exclusively by being relatively faster than other algos. Once typical response times are milliseconds, microsecond differences start to become important. It is not absolute speed that matters, but relative speed. As long as there are frictions in hardware or software there will always be benefits to being relatively faster regardless of the overall level of speed at which algos operate. Indeed, High Frequency Trading may be a misnomer – it should be Higher Frequency Trading.

Thus, while absolute speeds may approach the speed of light, as long as the current trading protocol prevails there will always be stiff competition to shave off ever-smaller time increments in overall execution speed. In Section 3 we argue that this competition is and will continue to be socially harmful.

2.2. How much is speed worth?

2.2.1 The private value of relative speed

The quest for speed is driven by a clear market incentive: In an electronic order book market, speed is worth a lot of money. The benefits of speed derive from three sources:

1. Aggressively exploiting or "picking off" stale passive orders.
2. Keeping passive orders from becoming stale or from being picked off by aggressive orders.
3. Obtaining a better position in order book queues than competitors with similar strategies and information.

---

[4] See http://www.algotechnologies.com/sitedata/Misc/AlgoM2_Latency_Comparison.pdf. See also AFM (2011).

[5] See Lorence (2011).

It should be noted that all these private benefits come from relative speed, so that the incentives for relative speed do not decrease with the absolute speed level of the market. Of course to the extent that increasing speed eventually harms markets, this will also be felt by private traders but this overall effect is privately negligible to each of them (it is a large overall effect divided across millions or billions of investors). So private incentives will cause market speed and therefore social costs and this could lead to indefinite increases in market speed as the driver is relative speed which will always be improvable with further investments.

Based on these observations, it is important to try to quantify the incentives for private speed as this can help understand the enormous magnitude of the divide between social and private incentives for market speed.

Most studies so far have focused on the profit potential involved in (1). HFT strategies of this type necessarily involve aggressive trading, in which high frequency traders pick off stale quotes or exploit momentary profit opportunities using market orders. The estimates for the profit potential of aggressive trading are surprisingly small, ranging from $1.5 billion to $21 billion per year worldwide.[6] We should note however we believe this number is a huge underestimate because it does not take into account the fleeting near-arbitrage opportunities generated from trading the same or similar instruments across different exchanges, which is known to be the bread-and-butter of many HFT firms. Unfortunately it is impossible for an academic or a regulator to estimate the profits available from fleeting near-arbitrage opportunities as this would require data synchronized at the location of trading servers and time-stamped within microsecond accuracy, which is available only to HFT firms. Even exchanges and regulators do not have access to this kind of data, let alone academics.

From a qualitative theoretical perspective, arguments by Cvitanic and Kirilenko (2010) suggest that the profits of high frequency “sniper” traders increase with the arrival frequency and variance of low frequency (human) traders, while Moallemi and Saglam (2010) suggest that the value of speed increases with asset volatility and decreases with the size of the bid-ask spread.

The profit potential for (2) is by construction the opposite of (1) – this can be thought of as an avoided loss rather than a profit. The losses to be avoided by keeping passive orders from becoming stale are by definition equal to the profit potential from picking off stale orders. This highlights the fact that the value of speed is not determined by the level of speed but rather by relative speed. In this sense, much HFT is the outcome of a needless arms race, in which aggressive “predators” continually attempt to prey on passive prey, and both of them evolve to be faster and faster without any clear social purpose. It also suggests that recent proposals to limit the value of speed for passive orders by limiting them to have a minimum duration will have an effect on overall market speed that is at least partially offset by increases in the value

[6] According to one early report (Martin, 2007), a 1-millisecond speed advantage can be worth $100 million a year to a major brokerage firm, though this estimate is not substantiated in any way. Kearns, Kulesza and Nevmyaka (2011), further develop the approach of Aldridge (2009) to make detailed calculations for an upper bound on the profitability of *aggressive* HFT on US equity markets. They arrive at estimates in the region of the single digit billions of dollars, which seem fairly low considering the extremely optimistic assumptions they make for the HF trader's forecasting ability, response times and transaction costs. They also report estimates for the profitability of HFT of several other authors which reach up to $21 billion per year with estimates to aggressive HFT of up to $3 billion per year. Arnuk and Saluzzi (2009) estimate profits from the "predatory" subset of HFT (defined as clearly harmful for markets) at $1.5-3 billion per year and Brogaard (2010) finds that HFT on the US equity market earns roughly $3 billion per year.

of relative speed for aggressive orders since more opportunities for picking off stale passive orders will become available.

Taking a different direction to the extant literature, we have recently considered the profit-making potential for the third source of profits, i.e. obtaining a better position in order book queues, and found that it is enormous (Skouras and Farmer, 2011). A rough estimate is that speed in limit order placement is worth at least $500 billion per year worldwide. This is at least one and possibly two orders of magnitude higher than the profits due to aggressive orders. The enormous size of these profits makes it clear why the evolutionary pressure driving speed is so intense.

Our argument hinges on analysis of the reasons that queue priority is advantageous. This is because (1) limit orders with queue priority trade more often and (2) orders with queue priority more often trade against small market orders, which have lower market impact, i.e. they move the market less in a disadvantageous direction. The latter effect is the more important of the two: Large market impact can undo the profits from passive quoting. The argument is described in more detail in a technical appendix to this document.

Based on data for global equities we use this result in Table 1 to obtain an estimate of profits from queue positioning in the order of tens of dollars per trade which implies that there are hundreds of billions of dollars to be earned from fast positioning in electronic order book queues.

spread (basis points)	3,5	4,68	12,05
quote size (US$)	691643	153215	1056675
number of trades per year (thousands)	18258756	1052208	5998188
Profits from queue priority per trade	33,62	9,96	176,85
Profits from queue priority annual ($ bn	**204,63**	**3,49**	**353,59**

Table 1: This table presents our calculations for profits to relative speed obtained through priority in queues of electronic limit order books. Our approach is discussed in detail in Section 2.2.1 and the Technical Appendix of this document. The spread and quote size values are based on our own calculations as 20 day averages ending on April 27 2011 for all stocks in all markets covered with valuations in the region US$64-128 billion. The number of trades were annualized based on values reported for May 2011 at www.world-exchanges.org/statistics/key-market-figures. As a rough approximation, we assume that one third of all trades occurred in these large cap stocks and ignore profits from smaller cap stocks. The β power-law exponent we used was -2.8 which we estimated on one year of tick trade data (2009) for large cap European stocks. The estimate was very close to -2.8 for each stock, for details see Skouras and Farmer (2011).

Certainly refinements to our calculations are possible but we believe this calculation is a powerful way to show that the value of relative speed is extremely high in electronic limit books and that this may be largely due to the popularity of the time priority rule, which is not a law of nature – barring political barriers due to vested interests, this is easy to change. We note that two top HFT firms, Getco and Quantlab Financial each spent around $300,000 on lobbying and $100,000 on political contributions in 2010 (Bowley, 2011).

Ta good approximation, the quest for speed is a zero-sum game with the obvious flavor of an arms race. Since it is this arms race that determines the level of absolute speed, it is obvious that society will overinvest relative to what would be an ideal social optimum. Indeed, if as we argue social welfare eventually decreases with market speed then this overinvestment can be very substantial. We address these issues in detail in the subsequent sections.

2.2.2 Empirical evidence on the effect of market speed

As an empirical matter the effect of market speed has been analyzed in the context of HFT and the evidence is somewhat mixed. On the negative side, the comprehensive analysis of Boehmer *et al* 2012 finds that HFT decreases liquidity and increases volatility. In earlier work, BMO Capital Markets (2009) argues that the effect of HFT on the Canadian Market has been detrimental and Zhang (2010) finds that HFT tends to increase volatility. Zhang and Powell (2011) suggest that HFT was responsible for the flash crash of 2010 and caused an adverse impact on market confidence that they believe has been large. Jarrow and Protter (2011) and Cartea and Penalva (2010) develop models in which HFT is generally harmful.

On the positive side, Jovanovic and Menkveld (2010) find that a HFT market maker in Dutch stocks leads to a reduction of spreads by 29% while Brogaard (2010) finds that in the US equity markets HFT contributes to price discovery and is generally beneficial.

In our view this mixed evidence is due to the fact that HFT is a very rich subset of strategies, some of which are beneficial for markets while others are not, with these effects being potentially dependent on other market conditions. It is misleading to lump all of HFT together and a detailed ecological approach is necessary to empirically evaluate the effect of HFT. This is impossible without data that identifies participants and a new body of research making use of this data. Farmer and Skouras (2011) develop this perspective and propose a research agenda in this direction.

One thing that is clear about HFT is that it has distributional consequences for the profits available to traders: Tradeworx (2010) describes one situation in which HFT can - and in their experience does - exploit market microstructure imperfections to the detriment of other traders. A QSG study (2009) reports that HFT strategies have been designed to exploit features in execution algos. One way to interpret these reports is that there is a predator-prey relationship between fast and slower traders.

Acknowledging the limitations current empirical evidence has for evaluating the effects of HFT, we turn to a discussion of the effect of market speed which is based on more theoretical reasoning from first principles.

3. Risk assessment

We now turn our attention to evaluating the costs associated with taking no measures to regulate market speed.

It is far from obvious that HFT is socially harmful. It is widely viewed by many as the outcome of an inexorable process of competition and technological innovation that in many other domains has been beneficial to markets (Jain and Johnson, 2008) and to society at large. Since HFT technology is available to anyone who is willing to make the investment in it, it seems subject to the usual competitive pressures and so markets should function properly to deliver reasonable outcomes. Furthermore, as discussed in the introduction, financial markets have always conferred an advantage to speed and considerable resources and ingenuity have always been spent to reduce latencies so it is possible to argue that HFT is not a qualitatively new phenomenon, even if it is more important now in terms of scale than it ever has been.

But it is difficult to find theoretical reasons why current speed levels might confer significant benefits to society. One argument might be that if prices reflect information sooner rather than later, then more of the information that influences the prospects of an investment is available at the time of a decision and therefore its risk is smaller. Taking as an example an investment in the shares of a biotech company developing a new technology, this effect would mean the company would be able to raise capital at a slightly lower cost. Society at large would benefit because of reduced uncertainty about the future. The problem with this argument is that any benefits of uncertainty reduction from receiving information milliseconds or microseconds earlier are surely negligible yet these are the speed levels at which markets operate. The opportunity to share risks of negative events that might occur within milliseconds seems an unlikely source of sizable private or social benefits.

A second argument might be based on the idea that one of the main benefits of markets come from the information they generate about what is the appropriate price of securities. If faster markets mean markets perform this role more effectively then social benefits could arise because of this improved transmission of information. However, while faster markets may cause prices to reflect true values sooner, there are no obvious reasons why they would reflect true values more accurately.

Finally, a third argument in favour of market speed stems from the idea that markets create “gains from trade”. Trade is on average beneficial to both parties, for example because it allows risks to be shared. Ergo obtaining these gains sooner must be beneficial. But at plausible discount rates, obtaining these benefits sooner by a few milliseconds surely has negligible consequences.

On the other hand, there are several reasons why HFT might be harmful. Farmer and Skouras (2011) discuss in some detail the reasons for which algorithmic trading (more broadly than HFT) can cause market instabilities and why it requires a carefully designed competition policy, so here we will focus only on reasons which are specific to HFT beyond the general problems with algorithmic trading not specific to speed.

1. **Speed as a perverse systemic outcome**: Although speed has always had value in markets, it is not in any clear sense “natural” or inevitable that it should have as much value as it does presently. We believe its value now is at historical highs based on calculations that place it in the order of hundreds of billions of Euros per year. Our calculations reveal that this is so high because of a combination of the increase in the volume of trading and a reduction in frictions such as costs which previously cancel any small edge that a small differential information speed could offer. Indeed, we believe the importance of speed emerged as a byproduct of the symbiotic relationship of exchanges and HFT which itself is caused by the fact that exchanges have volume based fees (e.g. they charge per transaction) and HFT is by its nature high volume since it involves exploiting very many small profit opportunities. Small profits per transaction are made but because volumes are very large the profits become significant. In this context, exchanges competing for the business of HFR firms have an incentive to adopt market designs that are attractive to HFT, such as time priority electronic order books. This allows HFT to extract profits from slower traders, part of which is extracted by the exchange in the form of fees. This explains why markets use volume rebates in their pricing models to attract high volume HFT and why the Deutsche Börse recently went at great lengths to get permission from its regulator to allow algorithmic trading to have lower prices than other types of trading. This is consistent with the evidence presented by Menkveld (2011), which suggests a symbiotic relationship between Chi-X and a single large HFT market maker; it also indicates there are strong path dependencies in

how the industry has evolved and why HFT has achieved its current levels of importance.

The profits from queue priority do not exist at all in “pro-rata” limit order book markets in which there is no time priority but instead orders at the same price execute against incoming aggressive orders in proportion to their relative size.[7] The fact that new market designs are emerging such as “dark pools” which are designed to exclude HFT suggests that the value of speed may be an artifact of coincidental systemic characteristics of today’s markets. De facto weak monitoring and regulation of HFT is another indication that institutions are set-up to favor HFT in ways that would not be considered acceptable for more mature types of trading. A key role has also been played by previous regulations which kept tick sizes at large sizes which in turn has led to higher value of time priority in order book queues.

2. **Speed creates tradeoffs that increase risk**: In order to increase speed, there is no doubt that algorithmic traders take short-cuts when it comes to implementing risk controls or in handling unusual situations. If speed is a priority, this necessarily limits the processing time that is needed for other tasks, such as risk control or “intelligent” information processing. Obtaining the sophisticated understanding needed to evaluate the meaning of information requires large programs that consume many clock cycles. Without such processing it is impossible to make intelligent decisions. Even though processing power gets continually faster, speed requirements also get faster, and the sophistication of information processing never improves – the algorithms just get faster.

   “Naked access” is a type of extremely fast market access that became popular (until it was eventually banned in the US) because it allowed HFT to bypass risk controls on the broker side. While adding safety for the client, broker and market, and having been universally required for many years, HFT firms managed to convince brokers to give them naked access to reduce the latency of their orders. The drawbacks of this are obvious.

   Controlling market speed is important in the same way that it is in NASCAR racing. Drivers have such strong incentives to go fast that they are almost forced to use technology that causes risks that will eventually kill them unless appropriate rules are imposed by a regulator. Imagine the outcome for example if racing cars had no constraints on how they could be designed or powered while drivers were allowed to dope themselves as they saw fit in order to optimize their performance.

   While it is possible that speed can be used for immediate hedging or other risk mitigating factors we doubt it is of much consequence whether this happens in nanoseconds or seconds.

3. **Speed can cause systemic instability and collapse**: HFT has dominated the market ecology but may disappear in abnormal conditions, which can cause a disruption in the ecology through a sudden reduction in diversity. Research in market ecology suggests this can be very dangerous for stability (see Farmer and Skouras, 2011). HFT can pull out of the market very easily as HFT firms usually have small inventories.

[7] This pro-rata approach is common in futures contracts, especially for interest rates.

There is one line that is clearly the fastest connecting the futures markets in Chicago with the Equity / ETF markets in New York. The highest frequency traders all use this same line and are responsible for a large fraction of the volume on both exchanges. If this line fails there would be a huge disruption in both markets that could easily spread globally and indeed, such a vulnerability could be attractive to terrorists. The broader point is that this vulnerability is caused by high frequency trading, while if relative speed did not confer such a huge advantage, network speed would not be a bottleneck.

It is widely acknowledged that the strategies used by HFT are highly correlated so HFT can lead to systemic risk. The heterodox explanation for the Flash-crash given by Nanex (2010) suggests it occurred because highly correlated automated market making algos caused a surge in activity, which caused an overwhelming data processing task, which triggered the algos to attempt to close positions more or less simultaneously in an extremely aggressive sell-off. This illustrates a problem with HFT whether or not the explanation is correct. Less pressure for speed would allow the algorithms to be more complex and more diverse, and therefore (likely) more stable.

4. **Speed causes instability in market location**: The increasing demand for market speed means there are huge benefits to markets locating in physical proximity to each other. This is already disruptive to the historical pattern of the existence of national exchanges and may well lead to all markets being concentrated in a single location. While this is not in itself harmful on a global level, it will be damaging for most current financial centers which - if trends continue - will eventually cease to exist.
5. **Speed as a barrier to entry**: Current HFT institutions have a huge incumbent advantage, with costly barriers to entry and an oligopolistic ownership structure. For example, Arnuk and Saluzzi (2009) cite a TABB report according to which expenditures on colocation and facilities for fast access amount to $1.8 billion per year (we assume this number refers only to expenses in US equity markets). In addition, markets are also spending huge sums and the NYSE alone is investing in facilities at a cost of $500 million. According to Price (2009), the cost of in-house solutions for competitive data feeds is of the order of $260,000 *per month* and a start-up cost of $270,000 per data center. Just to begin recording the kind of data that is necessary as a first step in designing a HFT strategy, a potential entrant would need several million per year. Furthermore, brokers typically require very large amounts of capital under management in order to provide competitive commission schedules (of the order of several tens of millions). The industry is very opaque in terms of information regarding fees and latencies of competing brokers and connectivity providers and it is a very difficult task to collect and compare such information across the large menu of available options.

One reason for the oligopolistic nature of speed is that a low frequency trader does not have access to the historical data with accurate, high resolution timestamps that are necessary to develop and backtest many types of HFT strategies. Furthermore, in our experience HFT firms demand ownership of any strategies that are developed using their data or infrastructure, making it very difficult for new start-ups to enter the HFT space.

While it is true that competition for speed has always been present in financial markets, it is also true that regulators have always been involved in trying to level the playing field as much as they could. This has become increasingly difficult because of speed itself which makes it much harder for regulators to even observe what is going on in markets (on this see also Section 4.3). The fact that there are other barriers to entry such as the availability of expensive research or information is no reason to not try to level the playing field where this is feasible, if there is strong evidence that eliminating such barriers is beneficial.

Indeed, there is evidence that policy to eliminate barriers to entry or to level the playing field for traders can have very positive effects. For example, Easley, Hendershott and Ramadorai (2009) study a natural experiment in which market access was changed to significantly equalize speed access of on and off-floor traders on NYSE and found that it led to an increase in prices of 3%, suggesting that this kind of issue can have a huge effect on the behavior of markets. Easley, O'Hara and Yang (2011) use a theoretical model to argue against practices which allow differential access to information (including speed advantages to HFT, an issue which seems related to the Net Neutrality debate relating to whether homogenous access to the Internet should be enforced by regulation - see e.g. Economides, 2008). That *equal* access to markets is viewed as a very important consideration is underlined by the fact that co-location providers make every effort to ensure that an identical service is provided to all co-location customers (e.g. the length of wires from market servers to all customers' co-located servers is the same regardless of where the servers happen to be located in the market data centre).

In sum, we acknowledge that barriers to entry are inevitable in markets, that speed is only one of them and there is not necessarily any foul play in how they emerge. However we believe that regulators should nevertheless try to contain these barriers to entry where practical as this is beneficial for markets. It may also be beneficial for market participants including those that build the barriers who may be subject to winners' curse (i.e. they spend so much to be the fastest that this offsets their advantage).

6. **Speed causes a wasteful arms race**: The analysis of section 3 leaves no doubt that speed competition has become an arms race with "Red Queen" characteristics, i.e. "It takes all the running you can do, to keep in the same place". Effectively this results in a transfer of wealth from investors and market participants to technology participants who are the "arms dealers" in this situation. Further improvements in latency technology are unlikely to improve the behavior of the markets but will perpetuate this transfer of wealth indefinitely if the arms race is not contained.

   A further disturbing feature of this arms race is that it may have a winner-takes-all outcome, i.e. if an institution invents a way to be faster than everybody else they can take all the opportunities known to be available to the very fastest algos (e.g. pure arbitrage opportunities). This will eventually cause competitors to exit the arms race, decreasing competitive pressure on the winner. Indeed, market lore suggests that some latency sensitive near-arbitrage opportunities are already entirely dominated by single institutions that have tailored their infrastructure especially to exploit the specific opportunity.

7. **Speed makes effective regulation much harder**: Speed itself causes regulators to have a hard time forming effective policies. The technological arms race between HFT firms means that their data analysis resources are far superior to those of regulators. Both regulators and academic researchers have incomparably smaller personnel and incentives to collect and analyze market data than HFT firms. As a

result the capability to analyze many issues of concern lags reality because markets change faster than the needed studies can be done. It is indicative that while the Flash Crash lasted half an hour, it took more than six months for the SEC team of experts to prepare its report. Slowing down this arms race would make it easier for regulators to understand markets and perform their functions. It is an extremely complex task for regulators to evaluate all the market innovations that speed leads to and whether they are harmful or beneficial, including flash orders, naked access and dark pools. When they do conclude they are harmful this is always with a lag during which harm has been done; this lag is the time period over which HFT innovators can expect to make profits from new trading technologies. As an example, we note that after several years since the appearance of flash orders, they are still allowed on some exchanges even though they involve front-running which is generally illegal.[8]

The potential sources of costs and benefits associated with speed are summarized in Table 2.

Sources of potential benefits	**Sources of potential costs**
Faster price discovery	Speed may be achieved at the expense of more operational risks
More accurate price discovery	Increase in probability of systemic instability & collapse
Gains from trade sooner rather than later	Increase in instability in market location
	Reduces competition because it is a barrier to entry
	Wasteful arms Race
	Regulation is harder

Table 2: Summary of sources of potential costs and benefits of speed. See Section 3 for a detailed discussion.

In sum there are many reasons to believe that speed is likely to have harmful effects at its current level. Furthermore, studies of the effect of market speed-ups in the range of seconds do not seem to find a significant improvement in market quality (e.g. Webb et al, 2007) so we believe that we have long reached the point where diminished returns to speed have kicked in. At the same time, we fully acknowledge the need for rigorous empirical analysis of each of the points we have made, a task which would no doubt involve tens of man-years of research. Unfortunately research is not as fast as markets and policy makers will need to make judgments based on the available qualitative evidence.

4. Options: our proposal for optimizing market speed levels

Recent radical regulatory proposals by prominent politicians (e.g. to ban short-selling or impose transaction taxes) are typically not grounded on empirical evidence. Viewed from an ecology perspective such proposals are risky as they would indiscriminately kill large and varied niches of trading species with potentially huge consequences for the entire ecology (see Farmer and Skouras, 2011). In the medium run, policymakers must collect the data, infrastructure and expertise that will make an evidence-based approach feasible.

[8] As a referee pointed out, flash orders are closely related to HFT and could be the subject of a separate regulatory impact appraisal.

Meanwhile, building on the detailed considerations discussed in the previous sections we propose a more subtle way to contain the benefits to speed which will thereby cause a reduction in speed itself without having much direct impact on trading strategies that are not predicated on speed. On a broad level, we suggest the following:

1. The value of speed must be reduced while otherwise affecting markets as little as possible.
2. Any policy must be global, i.e. implemented at the G-20 level to avoid regulatory arbitrage. Care must be taken to ensure that internal crossing networks or over the counter markets are not used to bypass regulations.

We have a well-defined proposal for how to achieve the first goal which we develop below. However, we emphasize that more work needs to be done in order to evaluate the proposal and work out many important details.

Our proposal is to require the following features for all trading mechanisms:

- **Pro rata:** Electronic order books should replace time priority rules with pro-rata rules. In modern electronic order book markets, orders can be sent at prices such that there is no immediate interest against which they might execute. Such orders are “passive limit orders” and there will often be several orders at the same price. Markets have rules regarding how such passive orders will execute in the event that there are incoming aggressive orders which can fill some but not all of them. In the usual “time priority” markets, passive orders at the same price are allocated against incoming orders on a first-come first-serve basis. By contrast, in pro rata markets, all passive orders get partially executed, each in proportion to its size. This will drastically reduce the advantage to processing information and sending orders before competitors. A cap on the size of orders will be required so that institutions with large capital do not overwhelm the queue. This will have minimal implications for aggressive (market) orders even though it reduce the private value of relative speed for passive (limit) orders which we argued in Section 2 is huge.

- **Frequent sealed bid auctions:** All market venues will be allowed to process events and disseminate information only once every second, within a few microseconds of a globally common clock reference. According to the terminology of Schwartz and Francioni (2006) these would be ‘sealed bid auctions’ because orders are not visible except at the times of auction trades. As discussed below, this time interval is short enough to allow for many opportunities for intraday price discovery and to contain various potential costs and uncertainties that arise from the lack of trading opportunities between auctions. At the same time, it is long enough to stop the technological arms race driving algorithmic trading. Economides and Schwartz (1995) were very influential in the adoption of call auctions side-by-side to continuous trading mechanisms, especially at the open and close and the empirical literature examining the impact of these auctions has found overwhelmingly positive effects. However, this does not mean that frequent sealed bid auctions are necessarily a good replacement for continuous trading mechanisms and indeed the reasons for which we suggest this mechanism have not been studied in the literature as far as we are aware. The mechanism has been used (at up to 10 second intervals and without the other elements we propose) on TAIFEX for Taiwan Stock Index Futures Markets quite successfully at least relative to open outcry continuous trading according to several studies (see e.g. Huang 2004). Evidence from Webb et al (2007) suggests market behavior is not very sensitive to the frequency of trades, and indeed in TAIFEX the frequency gradually increased from once every two minutes to once every 10 seconds without a clear impact on market quality. Cheng, Kang and Fu (2008) find that the transition from this mechanism to continuous trading had mild consequences that depend on the liquidity of the underlying instrument.

- **Randomization/exponential durations:** The one second duration between event processing should in fact be random, with the probability that they will occur at any given instance being unpredictable; in other words, it is the average duration that should be one second. The random duration (common across all markets) would also eliminate the possibility that exogenous events such as macroeconomic / earnings announcements could create predictable opportunities for institutions that know they are faster (for example if the announcement takes places 1 millisecond before the time at which it is known with certainty that a market event will happen, fast traders will know they can exploit slower traders who will not be able to react in time). By making the timing of the market event uncertain, the profits available to fast traders will also go down as they will not be guaranteed a "last-mover advantage". Randomization is commonly used in opening and closing auctions especially in European exchanges but we are not aware of any studies specifically of the effect of randomization.

This shift from continuous-double auction electronic order book markets with time priority to markets with very frequent pro-rata call auctions will level the playing field so that it is accessible to anyone with an infrastructure equivalent to a Bloomberg terminal (a cost of around $30k per year). We mention this just as an example, since a Bloomberg terminal coupled with a high level programming environment can support the development of trading algorithms with around one second latency. Arguably, the chunking may be even more effective if it is less frequent, e.g. once every 10 seconds to allow humans to compete with algos in many activities where this is now impossible, but we believe it would be preferable to begin with a less disruptive intervention.

This is a drastic change given that as we have seen, cutting edge players now spend tens of millions per year in infrastructure. Our proposal would simultaneously end the arms race for speed and reduce the barriers to entry in the algorithmic trading space. It has the advantage that it allows markets to remain located at any distance from each other across the globe. Meanwhile, at one event per second markets can still generate almost 30,000 trades per (eight hour) day to facilitate price discovery through the trading process. Based on our review of the literature, there is no clear evidence that this approach or the pro-rata rule would have an adverse effect on liquidity which a critic might want to counterpoise to the obvious beneficial effects of our proposal (see section 5 for more details).

Any policy should ideally be global, i.e. implemented at the G-20 level to avoid regulatory arbitrage. Care must be taken to ensure that internal crossing networks or over the counter markets are not used to bypass regulations. However, we think our proposal below may work even if imposed unilaterally. Our reason for this is that it offers a trading environment that most market participants would welcome, and gives no selective advantage to either aggressive orders or passive orders. It is an environment that would make participants feel safe. Some evidence in this direction is that increasing volume in the last few years is being traded in 'dark pools' and other exchange mechanisms designed to exclude high frequency traders. However, we emphasize that more work needs to be done in order to evaluate our proposal broadly and work out many important details including whether a unilateral implementation would be successful.

Most importantly, this proposal would slow the market down and refocus the effort of market participants on more productive activities. Algorithm designers would be able to devote substantial clock cycles to risk control and other activities that allow for intelligent and stable price formation.

5. Costs, risks and benefits of replacing the continuous double auction with frequent pro rata sealed bid randomized auctions

Summarizing our discussion, we expect there may be huge benefits if our proposal to switch to frequent pro rata sealed bid randomized auctions were adopted because this would bring the private and social value of speed much closer to each other. Our argument has two elements (1) there are benefits to markets and society more broadly for reducing the absolute levels of market speed which have been described in detail in the Risk Assessment of Section 3; (2) market speed can be regulated quite accurately using the frequent pro rata sealed bid randomized auctions described above. We propose a particular speed level that seems reasonable (1 Hertz) but this level can easily be fine tuned after appropriate experimentation if our proposed market structure were adopted. Indeed, it might be desirable to set speed as a function of trading volume at a market or instrument level. However, the fastest market should still be at around the level of 1 Hertz to ensure that equal access is available to all traders with basic infrastructure such as a Bloomberg terminal.

The potential costs from this switch emerge from the difficulties in its implementation and from some frictions it creates in market interactions. The hardest part to implement is no doubt the synchronized randomization across all market venues and its global implementation. However, relaxing the synchronization requirement and the fully global implementation would still likely yield significant benefits. There are many variants of our basic proposal which might be simpler to implement, for example a random latency imposed by the exchange in the processing of any instruction (order submission/modification/cancellation) could potentially be calibrated to decrease the value of *relative* speed which we argued in Section 2 is the ultimate source of the problem.

In terms of frictions, a source of potential costs might come from reduced market quality under the new mechanism. This may occur for example because of an increase in execution uncertainty and execution price uncertainty for high frequency traders. However, we note that low frequency participants are already subject to these uncertainties and may in fact find it easier to monitor slower markets under our proposal. Another reason our proposal might have negative implications is that by reducing the profitability of fast liquidity provision it may reduce the positive externalities emerging from private incentives to provide liquidity.

The potential impact on market quality can only be inferred very imperfectly through empirical evidence from markets with similarities with the design we propose. Indeed, there is some evidence that pro-rata markets are less liquid than continuous markets (Lepone and Yang 2012). However, it is difficult to draw general conclusions from specific studies as the details are likely to matter and our proposal combines several elements. For example, it is quite plausible that continuous trading is preferable to daily call auctions (as reported e.g. by Amihud, Mendelson and Lauterbach 1997) even though it may be very similar to call auctions every ten seconds (e.g. as in Cheng, Kang and Fu, 2008). A more targeted study of our proposal would be welcome but we do not see a mechanism through which it could do significant harm to market quality. No doubt however the magnitude of the gains will depend on the details as has been repeatedly observed in the study of market design (see e.g. Kremer and Nyborg, 2006).

Of course there will be some market participants who will lose from adopting the proposed auction mechanism, especially stakeholders in the private value of speed. These include trading operations as well as market venues the competitive edge comes from their own speed. Some market venues in fact may become redundant in slower markets. For example, from discussions with market participants we understand that an important reason for which “dark pool” markets have become popular is because such markets decrease the advantage of high frequency traders. To the extent this is true, if our proposal were to be successful it would undermine the commercial viability of dark pools and related markets. This would seem of relatively little import other than to the stakeholders in the companies that own these market mechanisms. However, as discussed these costs are outweighed by the benefits to other sectors.

One may reasonably ask why market forces have not delivered a trading mechanism like the one we suggest if it is as effective as we believe. We have already given a preliminary answer to this in our discussion of “speed as a perverse systemic outcome” in Section 3.3. Furthermore, we believe such a mechanism is unlikely to evolve without regulatory intervention because competition among markets has essentially become a competition to attract the handful of institutions that provide liquidity and these institutions use high frequency technology. It is therefore very difficult for a “slow” market to threaten incumbent fast markets, even though this might be a more efficient market structure and could have been the observed outcome were it not for path dependency and various historical accidents including previous regulations. For the same reason, it seems unlikely that markets will choose to slow themselves down voluntarily.

6. Future: evolution of costs, risks and benefits in the next ten years

It is hard to imagine what markets will look like ten years from now, whether our proposals are accepted or not. What we can say with confidence is that unless regulators take a sophisticated approach to reducing the private value of speed, billions of dollars will be wasted with the goal of achieving minute relative speed advantages that confer private advantage but little social value. In the next decade or century, what “minute” means will change with technology but unless regulators intervene, the basic nature of the speed race is likely to remain the same.

7. Summary

We have argued that the private value of speed is much larger than has previously been appreciated and that this is the main reason we are seeing such an explosion in high frequency trading as technology has improved and transaction costs have gone down. We also argue that while difficult to measure, the social value of speed is very negative - in particular many individual investors are harmed by market speed, each by a relatively small amount, but in aggregate this harm is much larger than the benefits from speed which is extracted by just a few institutions. We present a proposal to reverse this situation by replacing continuous markets with frequent randomized pro rata sealed bid call auctions. This proposal has the advantage that it that would leave other aspects of markets largely unaffected.

Before adopting our proposal it would be advisable to conduct an in-depth study of its likely implications. We expect such a study would involve:

- Quantifying the private and social value of speed, refining the approach we outline here.

- A better understanding of the pros and cons of pro-rata markets for liquidity provision and price formation. It should be noted that little is known about the properties of pro-rata markets compared to time priority markets. The studies that do exist, seem to suggest small differences (Frino, Holl and Jarnevic 2000). While it is understood that the incentives involved can be quite different (Field and Large, 2008) the impact of these differences has not been investigated. Many academic papers developing models for order books in fact assume pro-rata execution even though they purport to model real world order books that have time priority (e.g. Glosten, 1994).

- A better understanding of the pros and cons of frequent call auctions relative to continuous trading. It is worth noting that such a mechanism has been used at the TAIFEX for trading the Taiwan Stock Exchange Stock Index Futures so could be analyzed in detail.

We believe the arguments in favor of the protocol we are proposing here are strong and we would be surprised if the further study we are advocating above found this approach to be unfavorable to the currently more widely used time priority-based protocol. The motivation for the study is rather to develop refinements and modifications, especially to account for constraints that regulators may have in what policy options are practically feasible. We also believe a thorough detailed analysis is important to counterbalance the resources that stakeholders in market speed will undoubtedly invest in opposing this proposal or any significant reform. The benefits of our proposed change are substantial, and we believe it strongly merits further attention.

8. Technical appendix: review of our argument for estimating the private value of speed

Consider the profits to speed for passive quoting in the context of a market making strategy. A market maker offers to sell at price x + s/2 to a trader sending an aggressive buy order and when executed, attempts to buy back at price x – s/2, where s is the spread and x is the midprice. (The midprice is the average of the bid and the ask price and the spread is their difference). The market maker runs the risk of an adverse midprice movement between the acts of selling and buying. Indeed, because of market impact, there is a systematic bias: A buy market order tends to drive x up; the larger the market order, the higher its impact. If the impact of the market order exceeds s/2 then the market maker may systematically be forced to buy back at a price higher than x – s/2, and should expect to take a loss.

Define a “small” market order as one with impact m < s/2, and a large market order as one with impact m > s/2. At any point in time there will be some collection of quotes at the best bid each with its own position in the priority queue. Any specific bid quote will execute against an aggressive incoming sell order that is large enough given its priority level. Given that it executes, the quote will be profitable on average if the probability of small aggressive orders is large relative to that of large orders which would have correspondingly large adverse market impact (i.e. would move x downwards). When an aggressive order arrives, an offer with queue priority is guaranteed to execute at least partially, whereas an offer deep in the queue executes only if the market order is so big that it clears out all the preceding offers. Thus the higher the queue priority, the higher the expected ratio of small orders to large orders, and the higher the expected profit.

Conversely, profits decrease with decreasing queue priority. Once queue priority becomes sufficiently low, then the offer becomes unprofitable due to the fact that it will never be hit by a small market order. There is therefore a marginal queue position Q where uninformed offers will on average just break even.

Under the assumption that the market is on average efficient, and that quote placement is subject to the same information as is available to all traders, the queue should be exactly Q. By making appropriate assumptions it is then possible to compute the average profit advantage to all the quotes with position better than Q. Under the assumption that market impact is linear (see Farmer, 2002, Huberman and Stanzl, 2004):

$$m_t(v) = a_t v$$

where m_t is the impact at time t, of an incoming market order of size v which will depend on a time-varying coefficient a_t.

This assumption considerably simplifies matters because it rules out the possibility that passive orders may be profitable over multiple periods even if they are not in single periods. For example, it can be shown that if market impact is convex (which we exclude), a passive order may have negative profits in a single-period setting but positive profits in a multiple-period setting.

Additionally, if we make another empirically plausible assumption that market orders are distributed according to a power law (see e.g. Lillo et al., 2003, Farmer et al. 2005 for an empirical justification):

$$f(v) = v^{\beta}, \beta < 0,$$

then the average profit for orders at prices better than the marginal queue position Q is (Skouras and Farmer, 2011)

$$-\frac{s_t}{2}\frac{1}{\beta+1}\left[\frac{Q_t}{2} - 1 + \frac{1}{2Q_t}\right].$$

Of course there are several criticisms that can be made to this calculation. For example, it ignores transaction costs and rests on simplifying assumptions about market impact, the distribution of the size of orders, of information across traders, of all passive orders being at an optimized state at each trade and the relations between all of the above. Most significantly, we obtain a measure of the value of speed when all information is the same across all traders in which case the value of speed is greatest since the fastest trader can effectively use all information before anyone else. In reality, the fastest trader can only use a very small subset of information, otherwise he would be slowed down by the task of processing the information which is one reason our estimate may be exaggerated. On the other hand, our calculation neglects levels of the book beyond the best bid and offer which also gain from time priority.

References

AFM, 2011, “High Frequency Trading: The application of advanced trading technology in the European marketplace”, Netherlands Authority for the Financial Markets (AFM).

Aldridge, I., 2009, “How profitable are high-frequency strategies?”, FINalternatives.Amihud, Yakov and Haim Mendelson and Beni Lauterbach, 1996, “Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange”, Journal of Financial Economics 17, 223-249.

Arnuk S. and J. Saluzzi, 2009, “Latency arbitrage: The real power behind predatory high frequency trading”, Themis Trading LLC White Paper

Boehmer, E., Fong, K., and J. Wu, 2012, “International Evidence on algorithmic trading”, working paper.

Bowley, G., 2011, “Fast traders, in spotlight, battle rules”, New York Times, July 17, 2011.

Brogaard, J.A., 2010, High frequency trading and its impact on market quality, working paper

BMO Capital Markets, 2009, The Impact of High Frequency Trading on the Canadian Market, July 22, 2009

Bhupathi, T., 2010, “Technology’s latest market manipulator? High Frequency trading: the strategies, tools, risks and responses”, North Carolina Journal of Law and Technology, 11(2).

Cartea A, and J. Penalva, 2010, “Where is the value in high frequency trading”, working paper

Cantillon, E. and P-L Yin, 2010, “Competition between Exchanges: A research agenda”, International Journal of Industrial Organization, forthcoming

Chaboud, A., Hjalmarsson E., Vega, C. and B. Chiquoine, 2009, “Rise of the Machines: Algorithmic trading in the foreign exchange market”, FRB International Finance Discussion Paper No. 980

Cheng, M.S., Kang, H.H. and Y-F Fu, 2008, “Trading activity and the effects of trading mechanism”, working paper.

Copeland T.E. and D. Galai, 1983, “Information effects on the bid-ask spread”, Jounal of Finance, XXXVIII(5).

Cvitanic, J. and A. Kirilenko, 2010, “High frequency traders and asset prices”, working paper

Easley D., Lopez de Prado, M.M. and M. O’Hara, 2010, “The microstructure of the flash crash: flow toxicity, liquidity crashes and the Probability of Informed Trading”, working paper

Easley, D., Hendershott, T. and T. Ramadorai, 2009, Levelling the trading field, working paper

Easley, D., M. O’Hara and L. Yang, 2011, “Differential access to price information in financial markets, working paper

Economides, N. and R.A. Schwartz, 1995, Journal of Portfolio Management, 21(3), 10-18.Farmer J.D. and S. Skouras, 2011, “An ecological perspective on the future of computer trading”, Foresight Driver Review DR6.

Evans, D.S., 2003, “The antitrust Economics of multi-sided platform markets”, Yale Journal on Regulation, 325.

Farmer, J. D., 2002, “Market Force, Ecology, and Evolution.” Ind. & Corp. Change 11 (5) 895-953.

Lillo, F., J. D. Farmer, and R. N. Mantegna, 2003, “Master Curve for Price-Impact Function.” Nature 42, 129-130.

Farmer, J. D., P. Patelli, and I. Zovko, 2005, **“**The Predictive Power of Zero Intelligence in Financial Markets.” Proceedings of the National Academy of Science USA 102(6) 2254-2259.

Farmer J.D. and S. Skouras, 2011, “An ecological perspective on the future of computer trading”, Foresight Driver Review DR6.

Farmer J.D. and S. Skouras, 2011b, “Economic Regulatory Impact Appraisal of Amendment 2.3.f and question 20 of the European Commission Public Consultation for the Review of MiFID”, working paper.

Field, J. and J. Large, 2008, “Pro-rata matching in one-tick futures markets”, working paper

Foucault, T. and A.J. Menkveld, 2008, “Competition for order flow and smart order routing systems”, Journal of Finance, LXIII(1)

Frino, A., Holl, A. and E. Jarnevic, 2000, “An empirical analysis of price and time priority and Pro Rata trade execution algorithms in screen-traded markets”, Journal of Derivatives, Summer, 41-47.

Glosten, L.R., 1994, “Is the electronic limit order book inevitable”, Journal of Finance XLIX(4)

Gurliacci, M., Jeria, D. and G. Sofianos, 2009, “When the going gets tough, the algos get going, Journal of Trading, Vol 4, 2, 34-45

Hendershott, T. and P.C. Moulton., 2010, “Automation, speed and stock market quality: The NYSE’s Hybrid”, working paper.

Huang, Y.C., 2004, “The market microstructure and relative performance of Taiwan stock index futures: a comparison of the Singapore exchange and the Taiwan futures exchange”, Journal of Financial Markets, 7, 335-350.

Huberman, G. and Stanzl, W., 2004. Price manipulation and quasi-arbitrage. Econometrica, 72, 1247- 1275.

Iati R. ,2009, “The Real Story of Trading Software Espionage”, Advanged Trading, July 2010.

Jain P.K. & W.F. Johnson, 2008, “Trading Technology and Stock Market Liquidity: A Global

Perspective”, in, Stock Market Liquidity, F-S. Lhabitant and G.N. Gregoriou, John Wiley and Sons Inc.

Jarrow R.A. and P. Protter, 2011, “A dysfunctional role of high frequency trading in electronic markets”, Johnson School Research Paper Series No. 08-2011

Jovanovic, B. and A.J. Menkveld, 2011, “Middlemen in Limit-order markets”, working paper.

Kearns, M., Kulesza, A. and Y. Nevmyvaka, 2010, Empirical Limitations on High Frequency Trading Profitability, working paper

Kremer I. and K.G. Nyborg, 2006, “Underpricing and market power in uniform price auctions”, Review of Financial Studies, 17(3).

Lepone, Andrew and Jin Young Yang, 2012, “The impact of a pro-rata algorithm on liquidity: Evidence from the NYSE LIFFE”, Hournal of Futures Markets, 32(7): 660-682.

Lorence, G. (2011), “Informatica, HP and Mellanox/Voltaire Benchmark Report: Ultra messaging accelerated across three supported interconnects”, Informatica white paper, January.

McGowan, M.J., 2010, “The rise of computerized high frequency trading: use and controversy”, Duke Law and Technology Review, 16.

Matheson, T., 2011, “Taxing Financial Transactions: Issues and Evidence”, IMF working paper 11/54

Menkveld, A, 2011, “High frequency trading and the New-Market makers”, working paper

Moallemi, C.C. and M. Saglam, 2010, “The cost of latency”, working paper.

Nanex, 2010, “May 6th 2010 Flash Crash Analysis: Final Conclusion”, October 14, 2010

Phillippon, T. and A. Reshef, 2008, “Wages and Human Capital in the U.S. Financial Industry: 1909-2006”, working paper

Price, T. 2009, Market Data Management: In-house vs outsourced solutions, Towergroup, February.

QSG, 2009, “QSG study proves higher trading costs incurred for VWAP algorithms vs Arrival Price algorithms, High Frequency Trading Contributing Factor”, November 24, 2009

Schwartz, R.A. and R. Francioni, 2006, “Call auction trading”, in *Encyclopedia of Finance*, Lee C.F. and Lee A.C. (eds), Springer.

Skouras S and J.D. Farmer, 2011, The value of queue priority in limit order books, working paper.

Tabb L, R. Iati and A. Sussman, 2009. “US Equity High Frequency Trading: Strategies, Sizing and Market Structure”.

Tradeworx, 2010, “Public commentary on SEC market structure concept release”, April 21.

Webb, R, Muthuswamy, J. and R. Segara, 2007, Journal of Futures Markets, 27)12), 1219-1243

Zhang, F, 2010, “High frequency trading, stock volatility and price discovery”, Yale university working paper

Zhang, F, and S.B. Powell, 2011, “The impact of high-frequency trading on markets”, CFA magazine, March-April.

© Crown copyright 2012
You may re-use this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence. Visit www.nationalarchives.gov.uk/doc/open-government-licence write to the information Policy Team, The National Archives, Kew, London Tw9 4DU, or email: psi@nationalarchives.gsi.gov.uk

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk

URN 12/1072