



Government
Office for
Science

 Foresight

Economic impact assessments on MiFID II policy measures related to computer trading in financial markets.

Working paper

Foresight, Government Office for Science

This working paper has been commissioned as part of the UK Government's
Foresight Project on The Future of Computer Trading in Financial Markets.
The views expressed are not those of the UK Government and do not represent its policies.

Introduction by Professor Sir John Beddington

This working paper presents important interim findings of the international Foresight project: The Future of Computer Trading in Financial Markets. In particular, it considers the costs, risks and benefits of six possible regulatory measures which are currently being considered within the European Union's Markets in Financial Instruments Directive 2 (MiFID II). It precedes the final project report which will be published later in 2012, and which will consider a broader set of issues surrounding computer-based trading (CBT) over the next ten years.

Algorithmic trading (AT) and high frequency trading (HFT) have grown rapidly in use in recent years. As such, they have also fuelled increases in complexity as well as new system dynamics, making markets ever harder to understand and to regulate. In particular, there is continuing controversy concerning the extent to which they improve or degrade the functioning of financial markets, and also influence market volatility and the risk of instabilities. For example, such trading has been implicated by some as a contributory factor in the May 6th 2010 Flash Crash.

For such reasons, computer-based trading is now attracting the close attention of policy makers and regulators worldwide.

However, the debate on high frequency and algorithmic trading has been hampered by the availability of evidence and analysis. This is of significant concern since regulation that is not soundly based risks being ineffective, or worse, could lead to unhelpful and unforeseen consequences. By drawing upon the available science and evidence from across the world, the Foresight project seeks to provide independent advice to policy makers. More specifically, this working paper has involved some 35 leading academics from nine countries and presents analysis that has been subject to independent peer review. As such, it does not represent the views of the UK or any other government. In view of the rapid pace of the MiFID II regulatory process, I have pleasure in making this paper freely available now, in advance of the full Foresight report.



Professor Sir John Beddington CMG, FRS
Chief Scientific Adviser to HM Government and
Head of the Government Office for Science

Foresight project team

Professor Sandy Thomas (Head of Foresight)

Derek Flynn (Deputy Head of Foresight)

Lucas Pedace (Project Leader – Senior Economist)

Alexander Burgerman (Project Manager – Assistant Economist)

Gary Cook (Project Manager)

Jorge Lazaro (Analyst)

Christopher Griffin (Higher Executive Officer)

Piers Davenport (Executive Officer)

Zubin Siganporia (Research Assistant)

Yasmin Hossain (Research Assistant)

Louise Pakseresht (Intern)

Jennifer Towers (Intern)

For further information about the project please visit:

<http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading>

Contents

Introduction by Professor Sir John Beddington	2
Key findings	6
1. Notification of algorithms	7
2. Circuit breakers	9
3. Minimum tick sizes	14
4. Obligations for market makers.....	17
5. Minimum resting times.....	21
6. Order-to-execution ratios.....	24
7. Key interactions	27
8. Conclusions	28

Economic impact assessments on MiFID II policy measures related to computer trading in financial markets.

Oliver Linton, Cambridge University

Maureen O'Hara, Cornell University

J.P. Zigrand, London School of Economics

Published: August 2012

Key findings

Computer trading has changed markets in fundamental ways, not the least of which is the speed at which trading now occurs. There are a variety of policies proposed to address this new world of trading with the goals of improving market performance and reducing the risks of market failure. These policies include notification of algorithms, circuit breakers, minimum tick size requirements, market maker obligations, minimum resting times and minimum order-to-execution ratios. The Foresight Project has commissioned a variety of studies to evaluate these policies, with a particular focus on their economic costs and benefits, implementation issues and empirical evidence on effectiveness. This working paper summarises those findings.

The key findings relating to the different policies are as follows, starting with those which were most strongly supported by the evidence:

- Overall, there is general support from the evidence for the use of circuit breakers, particularly for those designed to limit periodic illiquidity induced by temporary imbalances in limit order books. Different markets may find different circuit breaker policies optimal, but in times of overall market stress there is a need for coordination of circuit breakers across markets.
- There is also support for a coherent tick size policy across similar markets. Given the diversity of trading markets in Europe, a uniform policy is unlikely to be optimal, but a coordinated policy across competing venues may limit excessive competition and incentivise limit order provision.
- The evidence offers less support for policies imposing market maker obligations. For less actively traded stocks, designated market makers have proven beneficial, albeit often expensive. For other securities, however, market maker obligations run into complications arising from the nature of high frequency market making across markets, which differs from traditional market making within markets. Many high frequency strategies post bids and offers across correlated contracts. A requirement to post a continuous bid-offer spread is not consistent with this strategy and, if binding, could force high frequency traders out of the business of liquidity provision. Voluntary programmes whereby liquidity supply is incentivised by the exchanges and/or the issuers can improve market quality.
- Similarly, minimum resting times, while conceptually attractive, can impinge upon hedging strategies which operate by placing orders across markets and expose liquidity providers to increased 'pick-off risk' if they are unable to cancel stale orders.
- The effectiveness of proposed measures to require notification of algorithms or minimum order-to-execution ratios are also not supported by the evidence. The proposed notification policy is too vague, and its implementation, even if feasible, would require excessive costs for both firms

and regulators. It is also doubtful that it would substantially reduce the risk of market instability due to errant algorithmic behaviour, although it may help regulators understand the way the trading strategy should work.

- An order-to-execution ratio is a blunt policy instrument to reduce excessive message traffic and cancellation rates. While it could potentially reduce undesirable manipulative trading strategies, beneficial strategies may also be curtailed. There is insufficient evidence to ascertain these effects, and so caution is warranted. Explicit fees charged by exchanges on excessive messaging and greater regulatory surveillance geared to detect manipulative trading practices may be more effective approaches to deal with these problems.

I. Notification of algorithms

I.1 The measure and its purpose

Algorithmic trading (AT) involves the use of computer programs to send orders to trading venues. Such algorithms now have widespread use among all classes of investors¹, and AT comprises the bulk of trading in equity, futures and options markets. Algorithmic trading is also fundamental to high frequency trading (HFT) strategies. A concern with algorithmic trading is that an errant algorithm could send thousands of orders in milliseconds to a market (or markets), resulting in major market upheaval. MiFID II Article 17(2) proposes that investment firms engaging in algorithmic trading must provide annually to the regulator a description of their algorithmic trading strategies, details of the trading parameters and limits, the key compliance and risk controls which are in place, and details of how its systems are tested. The purpose of this measure is to ensure that algorithms used in trading are subject to proper risk controls and oversight.

I.2 Benefits

If descriptions were able to prevent unsound algorithms from operating in live markets, then this would be a measure contributing to the maintenance of orderly markets. Requiring firms to have demonstrated risk controls in place might make aberrant algorithms less likely to occur. If regulators require increased testing at the firm level for algorithms it suspects are flawed, fewer events affecting market liquidity due to malfunctioning algorithms might occur.

One additional benefit is that regulators will have to acquire greater technical sophistication to understand and evaluate the algorithms being used in trading, which would improve their ability to investigate abusive practices. However, this would require substantial increases in personnel and greater investments in technology.

¹ The survey SR1 commissioned by the Foresight project found that algorithmic trading is used by 95% of asset managers, 100% of insurers and 50% of pension funds surveyed.

1.3 Costs and risks

There are substantial costs connected with meeting notification requirements in general, and particularly as currently stated in MiFID II Article 17(2). Cliff (EIA16)² argues that just providing a full description of an algorithmic trading strategy requires not only all the programs that have been written to implement it, but also the full details of the code libraries used, as well as the software tools involved. Moreover, these descriptions must include the actual computations required, the algorithms that affect the computations, and full details of how the algorithms are implemented. Providing information on the other aspects of the proposed regulation would be similarly complex. Regulators, in turn, would then have to analyse this material, and determine what risk, if any, this algorithm poses to the market. This would require substantial expertise at the regulator level with complex computing systems and analysis. An additional risk to consider is that algorithms are updated frequently, meaning that annual reviews will be ineffective in actually capturing the risk facing the markets.

Cliff estimates that it could run to approximately €1bn a year if the descriptions were in fact carefully read. Alternatively, the cost of MiFID II 17(2) could be dramatically lowered by simply having firms provide documents to the regulator that are filed but not really analysed. In this case, however, it is hard to see how this activity can actually address the potential risk of algorithmic disruptions to the market.

Care must be taken not to infer that this measure would dramatically reduce systemic risks, and that agents would consequently take larger risks than they otherwise would have. The reason systemic risk may not be reduced significantly even if algorithms were carefully analysed by regulators is that much of the risk arises from the nonlinear interactions of many algorithms. Different algorithms may be present in the markets at different times; setting up what could be infinite combinations of algorithms to consider for regulatory review. Furthermore, even if a 'wind-tunnel'³ for testing algorithmic interaction were constructed, it would not capture all of the systemic risks if algorithms learn and rewrite themselves dynamically over time.

1.4 Evidence

There is very little empirical evidence on the cost or benefits of algorithm notification. Cliff (EIA16) provides some estimates of cost, but these depend greatly upon how the notification requirements are implemented and on how much the required analysis is split between firms and the regulator. There is

² Throughout this document EIA refers to economic impact assessment studies commissioned by the lead expert group. These can be found in the project's webpage: <http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading/working-paper>

³ What is meant here is a simulated market that can in principle test algorithmic interaction much as a wind tunnel is used for testing aircraft designs.

to our knowledge no costs estimate of risk controls for algorithms at the firm level.

1.5 Conclusions

The desirability of understanding algorithmic trading strategies and their impact on the market is laudable but achieving this through notification requirements of the type currently envisioned in MiFID II may not be feasible given the complexity of algorithms and their interactions in the market.

2. Circuit breakers

2.1 The measure and its purpose

Markets have always been subject to episodic price instability, but computerised trading combined with ultra low latency creates increased potential for such instability to occur. This, in turn, has increased interest in the role and usage of circuit breakers. Circuit breakers are mechanisms for limiting or halting trading on exchanges. Their purpose is to reduce the risk of a market collapse induced by a sequence of cascading trades. Such trading could arise from pre-specified, price-linked orders (such as programme trades or trade algorithms that sell more when prices fall), or from self-fulfilling expectations of falling prices inducing further selling. Traditionally, circuit breakers were triggered by large price movements, and hence represented ex post reactions to excessive price volatility in the market. More recently, the advent of HFT taking place at millisecond speeds has resulted in a new generation of circuit breakers which work on an *ex ante* basis (i.e. halting trading before accumulated orders are executed).

Circuit breakers can take many forms: halting trading in single stocks or entire markets; setting limits on the maximum rises or falls of prices in a trading period (limit-up and limit-down rules); restrictions on one trading venue or across multiple venues. The London Stock Exchange (LSE), for example, operates a stock-by-stock circuit breaker that, when triggered, switches trading to an auction mode in which an indicative price is continuously posted while orders are accumulated on either side. After some time, the auction is terminated and continuous trading resumes. The trigger points are in several bands depending on the capitalisation and price level of the stock, the recent transaction history and the most recent opening prices.

However, circuit breakers are no panacea. Price discovery is a natural feature of markets, and bad news can induce (sometimes large) price drops to new efficient values. Halting markets can interfere with this natural process, and may simply postpone the inevitable. For example, the October 1987 crash in the US was subsequently followed around the world, with big price drops in the UK and other markets. The Hong Kong stock market was closed on the Monday after the US markets had started to fall and stayed closed for a week. When it did open, it suffered an instant decline of 30%. On the other hand, the May 2010 New York Flash Crash was effectively ended by a circuit breaker which allowed liquidity to re-accumulate as buyers returned to the market and the newly balanced market to resume. Thus,

circuit breakers, while well suited to dealing with instability caused by temporary shortages of buyers or sellers, are not appropriate for all causes of market volatility and cannot forestall revaluations that are unavoidable.

The policy issue is whether the existing self-regulatory, uncoordinated approach (in Europe) to price instability can be improved.

2.2 Benefits

If the price processes were driven purely by rational valuations of fundamentals, then a trading halt impairs the process of valuation and prevents the public from receiving accurate and up to date information. But in today's high frequency electronic markets, liquidity is uncertain, and prices can be affected (at least temporarily) by factors such as imbalances in the book of orders, fat finger trading errors, and errant algorithms engaged in a mechanical feedback loop. Circuit breakers and trading halts may be beneficial for dealing with such imperfections in trading processes. We describe the benefits of circuit breakers in the following three subsections.

2.2.1 Cooling-off period

Many modern markets function as computerised limit order books with continuous trading and replenishment of orders. Even if the daily trading volume is large, the displayed depth of the market at any moment may be relatively small. A large purchase order arriving unexpectedly, for example, can cause a temporary imbalance until more sellers come forward. Circuit breakers provide a respite that prevents mechanical selling at any price, allows the market to understand what is happening and gives counter-parties time to enter, thereby reducing the order imbalance.

In a fast-moving market, losses to positions bought using credit can build up quickly, leading the brokers who have provided credit to ask for additional collateral. With very fast moving markets, these margin calls cannot be satisfied quickly enough and broker confidence may suffer. A cooling-off period allows the traders to raise the collateral, reducing the risk that they fail. It also reduces market risk because brokers will not be forced to liquidate clients' positions which would then put additional selling pressure on the market and create a vicious feedback loop.

Circuit breakers can also be invoked for purely technical reasons to prevent peak overload bottlenecks at the exchanges' servers which could lead to erroneous pricing and execution.

2.2.2 Uncertainty resolution

Volatility is a natural part of markets, but unexplained volatility can cause traders to fear the worst and lead to massive selling. Because high frequency markets move so fast, it may be impossible for traders to evaluate what is causing a large price movement. There is now a large and growing literature showing that uncertainty reduces participation in markets, which can manifest in massive selling for those in the market and a reluctance to participate by those not already there. Either outcome is undesirable, so

mechanisms such as circuit breakers that can allow time for uncertainty resolution can be beneficial.

2.2.3 Investor protection

A market that is fast moving for reasons other than the inflow of fundamental news not only can create systemic risk but can also penalise traditional investors who do not have the resources to monitor markets continuously. Trading halts in response to non-fundamental swings may offer a means of preventing uninformed retail investors losing out to traders who continuously monitor markets. This could bolster the confidence of investors in the integrity of markets, and remove or ameliorate concerns that small investors can be taken advantage of by manipulative trend-generating strategies⁴. A final point is that circuit breakers enjoy widespread support from industry as shown in the survey SR1⁵. They are seen as a prudent mechanism to enforce orderly, fair and efficient markets.

2.3 Costs and risks

The obvious cost is that during a market halt, traders are prevented from completing mutually beneficial trades. An empirically documented effect of circuit breakers is the so-called 'magnet effect' whereby traders rush to carry out trades when a halt becomes imminent, accelerating the price change process and forcing trading to be halted sooner or moving a price more than it otherwise would have. Subrahmanyam (EIA4) recommends that there should be some unpredictability about when circuit breakers are triggered, a logic that may explain Deutsche Börse's decision not to publish the trigger points for its circuit breakers.

Similarly, a trading halt that slows down the fundamental price discovery process may create additional uncertainty. This can undermine confidence and increase bid-ask spreads when markets reopen if the reason for the halt is not credible. More generally, if circuit breakers are implemented across trading venues and coordinated so that a trading halt for a stock on one venue triggers a trading halt for that stock on all other venues, then a halt on a minor trading facility can trigger an unexplained halt on the main venue. In fact, experimental evidence suggests that completely random halts can create uncertainty because traders have the time to make up rumours or to exaggerate existing ones⁶.

⁴ Kim, Kenneth A. and Park, Jungsoo, Why Do Price Limits Exist in Stock Markets? A Manipulation-Based Explanation. *European Financial Management*, Vol. 16, Issue 2, pp. 296-318, March 2010. Available at SSRN: <http://ssrn.com/abstract=1560166> or <http://dx.doi.org/10.1111/j.1468-036X.2008.00456.x>

⁵ SR1 A survey of end-users. <http://www.bis.gov.uk/assets/foresight/docs/computer-trading/12-679-end-user-perspectives-on-computerised-trading.pdf>

Throughout this document SR refers to surveys commissioned by the lead expert group. These can be found in the project's webpage: <http://www.bis.gov.uk/foresight/our-work/projects/current-projects/computer-trading/working-paper>

⁶ See Ackert (EIA9) for experimental evidence on circuit breakers.

An additional detrimental element associated with circuit breakers is the inability for market makers to offload large positions quickly when trading is halted. This would have to be factored into their risk management system and might make them less willing to buy, reducing liquidity.

Similarly, there is an issue with respect to cross-asset, cross-venue trade. Suppose a trader is hedging a derivative or is engaged in arbitrage across securities. If trading in the stock is halted, the trader is suddenly no longer hedged and may suffer losses as a result. If this happens frequently, different markets may become less integrated, market efficiency suffers, and the possible loss of confidence in pricing accuracy can lead to feedback loops. If investors suddenly lose confidence in the price of a security (because trading has stopped) and they are invested in other assets in perhaps a carefully designed portfolio, they may decide to sell many other assets because their ability to control the total risk they face is compromised. This may lead to a chain of negative events across many securities.

A mandatory, market-wide circuit breaker has not existed in the UK, although in extreme circumstances LSE declares a 'fast market' and changes certain trading rules (for example, widening individual stock circuit breakers and relaxing market maker obligations). Due to the nature of HFT, much more trading now involves the futures market, exchange-traded funds (ETFs), contracts for difference (CFDs) and spread betting⁷. The US Flash Crash was triggered in the S&P500 E-mini futures market, then went to the ETFs on the index, and finally affected the equity market itself.

This raises the problem of how to implement circuit breakers across market venues. In the US Flash Crash, attempts by the New York Stock Exchange (NYSE) to slow trading were completely ineffective due to the ability of traders to use other venues which were not affected. For a uniform circuit breaker to be effective, it would have to close all markets for a single stock or series of stocks. This would require coordination between and across exchanges, but different exchanges have different trading rules and different trading practices. Moreover, the ability of an exchange to determine its own rules for handling volatility can be viewed as an important dimension of its risk management practices. Regulators would have to consider whether it is desirable for an erroneous trade on, for example, Chi-X to shut down the LSE, or whether the current system whereby only the primary market determines shutdowns is preferable⁸.

A hybrid system allowing regulatory override of individual exchange trading halt decisions might provide a mechanism to deal with market-wide disturbances.

⁷ In the UK, CFD's are often used as an indirect way of trading equities but avoiding the Stamp Duty on equity trading, since CFD's are not covered by this tax.

⁸ The costs and benefits of coordinated trading halts are further analysed in EIA20

2.4 Evidence

There is a sizeable literature on the impact of circuit breakers on markets, but overall the empirical results are mixed. Many authors find negative effects of circuit breakers, while others find no effects or small positive effects. However, it is difficult to analyse what would have happened had the market not been halted, and so with a few exceptions these findings are not statistically robust. A more important problem is that high frequency markets are very different from the markets analysed in previous research. There is, as yet, little academic research on the role of circuit breakers in high frequency market settings. In particular, many academic studies are on smaller international markets and are largely concerned with a simple type of breaker rather than the more complex models employed for instance by the LSE⁹. Furthermore, most analyses have almost exclusively focused on transaction prices in the immediately affected market rather than the order book or the spill-over effects on other securities and trading venues.

There is some evidence from the US Flash Crash on the effectiveness of modern circuit breakers. The end of the crash is generally attributed to the imposition of the Chicago Mercantile Exchange's 'Smart Logic', a circuit breaker that halts trading when the accumulated imbalance of pending orders, if executed, would result in the price falling beyond a pre-specified limit. This forward-looking circuit breaker differs from the variety generally employed in most markets which deal with issues once they arise, and may provide a template for the design of market-wide circuit breakers.

There is also some evidence on the use of single stock circuit breakers. The LSE describes how on an average day there are 30-40 trading suspensions, whereas in the first two weeks of August 2011 (when there was a great deal of volatility), this shot up to about 170 suspensions per day. Despite the high number of suspensions, large volume and wide market swings, trading was generally 'orderly' in their view¹⁰. Of course, there are other market-wide costs associated with each stoppage and it is not clear whether these costs are fully taken into account and appropriately balanced against the benefits.

2.5 Conclusions

Circuit breakers have a role to play in high frequency markets, and they are found in virtually all major exchanges. Because of the inter-connected nature of markets, however, there may be need for coordination across exchanges, and this provides a mandate for regulatory involvement. New types of circuit breakers triggered before problems emerge rather than after they have emerged may be particularly effective in dealing with periodic illiquidity.

⁹ A notable exception is a study by Abad and Pascual (2007) who investigate the type of circuit breaker used on the London Stock Exchange and the Spanish Stock Exchange.

¹⁰ The London Stock Exchange Group plc Response to ESMA Consultation: "Guidelines on systems and controls in a highly automated trading environment for trading platforms, investment firms and competent authorities" ESMA/2011/224, which is available at <http://www.londonstockexchange.com/about-the-exchange/regulatory/lsegresponsetoesmaconsultationonsystemsandcontrols.pdf> describes this.

3. Minimum tick sizes

3.1 The measure and its purpose

The minimum tick size is the smallest allowable increment between quoted prices in a market. Tick sizes have important implications for both transaction costs and liquidity provision. The transaction cost effect is straightforward: a larger tick size, if binding, increases trading costs by widening the spread between bid and offer prices. The liquidity effect arises because the tick size determines how easy it is for another trader to 'step ahead' of an existing limit order. In markets with a standard price-time priority rule, an order placed first executes ahead of one placed later unless the later order is posted at a 'better' price. Small tick sizes make it easier for traders to post that better price, so smaller tick sizes push up the cost for traders who put trades on the market and provide liquidity. The challenge for markets and regulators is to choose the optimal tick size to balance these liquidity and transaction cost effects.

There are important differences in minimum tick size policy between the US and Europe. In the US, Regulation NMS requires that in all 'lit' venues (exchanges and large Alternative Trading Systems) stocks over \$1 are quoted with a minimum tick size of one cent, and sub-penny pricing is prohibited. In Europe, there is no mandated tick size and local exchanges are free to set their own tick policy. As a result, there is generally a range of tick sizes depending on the price level and, in the case of LSE, on the market capitalisation. A European stock may trade on different public venues under different tick size regimes, whereas in the US, such differences can only currently happen in dark venues. Historically, the trend has been towards smaller tick sizes since US trading in 'eighths' (12.5 cents) yielded to decimalisation in 2000. Now, active stocks in the US typically trade at one cent spreads, leading to concerns that a one cent minimum tick may be too large. In Europe, spreads at minimum levels are not as common, suggesting that the tick rules are not as binding on market behaviour. The policy issue is whether it would benefit market quality to mandate a minimum.

3.2 Benefits

Choosing an optimal minimum tick size for a given stock and a given market environment would have a number of benefits¹¹. Originally, a uniform tick size rule was chosen to minimise transaction costs for firms and traders. With a limited number of possible prices, the technological complexity of trading was reduced and also the cost. While this is no longer needed given the state of trading technology, it is still the case that a well chosen minimum tick size framework can reduce the need for firms to split or reverse split their stock in order to influence the relative tick size. This latter reason speaks to the benefits of a non-uniform tick policy as found in Europe.

¹¹ For further discussion on the optimal tick size and its modelling see EIA7.

A well chosen minimum tick size can prevent market operators such as exchanges, market makers or high frequency traders making excessive profits at the expense of the final users. Chordia (EIA6) argues that a too high minimum tick in the US led to payment for order flow and internalisation as competitive responses to the excessive profits accruing to liquidity providers in large stocks. To the extent that trading venues compete with each other in a world with HFT and maker-taker¹² pricing, exchanges may have an incentive to try to undercut each other's tick sizes to attract volume and hence fees. A coherent overall minimum tick size policy (such as the one agreed to by Federation of European Securities Exchanges (FESE) members, discussed below) that applies to all trading venues could prevent a welfare destroying race to the bottom where competitive pressures would otherwise lead to a tick size below the optimal range.

A higher minimum limit for tick sizes can also provide greater incentives to post limit orders and thereby create a deeper, more liquid market. Because the costs of jumping the queue are higher with larger minimum spreads, there is less queue-jumping and so a generally more stable book. The reduced frequency of limit order book updating means less data and lower costs related to data. Whether such changes would reduce the arms race of trading firms investing in ever faster trading systems is debatable. A larger tick reduces the incentives to undercut, but it still remains the case that the first trader to post gets the order. This would imply that high frequency traders will generally dominate the order book for active stocks regardless of tick sizes.

Although Europe does not as yet have a mandatory tick size, there have been attempts by the industry through its organisation FESE to harmonize and simplify the tick size regimes across their members. The process has arrived at some agreement, but there is no binding legal framework to enforce it and it may not prove sustainable over the long term in the presence of new entrants.

3.3 Costs and risks

Setting too large a minimum tick size results in large bid-ask spreads favouring market makers. If prices must be set five pence apart, for example, then the bid-ask spread can also be no smaller than five pence. Except in special cases, a spread of one tick size cannot be undercut because prices must be quoted in increments of the tick size. A large tick size also reduces the chance of a price improvement. A price improvement occurs in any auction when a higher price is offered for the good. If the only bids accepted are in £1000 increments it is much harder to improve the price than if you are allowed to raise your offer by £1.

¹² Maker-taker pricing refers to the practice in many exchanges and trading platforms of paying a small rebate to executed orders that were placed as passive limit orders (the liquidity makers) and charging a fee to active orders that hit existing limit orders (the liquidity takers).

An excessive tick size on one market may induce activity to migrate towards venues with lower ticks. In the US, this is one reason given for the growth of trading in dark venues, and in Europe such tick size competition might be expected to induce orders to flow towards dark or Systematic Internaliser (SI) venues that are not subject to the rules for multilateral trading facilities. This migration is most likely to happen with low-priced stocks where the bid-offer spread is large as a proportion of the stock's value. Migration can be mitigated through additional measures, such as coordinated rules across trading venues and SIs, or via mechanisms such as the Retail Liquidity Providers (RLPs) proposed by the NYSE and recently approved by the SEC¹³. These RLPs would be allowed to quote on sub-penny ticks provided their quotes are hidden and can only be accessed by retail order flow.

Setting too small a minimum tick size comes with costs as well. In markets with price-time-priority, passive traders submit limit orders and thereby give away a free option to market participants. These passive traders expect to cover these losses through the spread. If tick sizes are very small, then new passive traders can come in and capture the incoming marketable orders by undercutting the current best bid or ask by a tick. Traders who provide liquidity by putting trades on the market will not be rewarded for having taken the risk of being picked off by traders with new information or through adverse selection. Too small a cost for jumping the queue, therefore, makes providing visible liquidity more expensive and leads to smaller and more ephemeral depth. It may also contribute to more cancellations and fleeting limit orders in the book as traders try to snipe in as late as possible.

3.4 Evidence

There is a large academic literature investigating the influence of tick sizes on market behaviour¹⁴. In general, the results from studies of a wide range of markets find that a reduction in tick sizes reduces spreads but also reduces depth. As a result, transactions costs for smaller retail investors tend to be lower, but the results are ambiguous for institutional investors whose trades are in sizes that may require greater depth. These empirical findings are consistent with the recent pilot programme implemented by FESE and subsequently analysed by BATS (2009)¹⁵.

Citicorp's recent reverse stock split underscores the effects that changing the relative tick can have on trading¹⁶. After Citi substituted one share for 10, trading in its stock shifted from alternative trading venues to exchanges; HFT activity increased; the trade size rose on alternative venues; volatility was

¹³The US Securities and Exchange Commission (SEC) approved the NYSE proposed pilot on July 3, 2012. See Release No. 34-67347; File Nos. SR-NYSE-2011-55; SR-NYSE Amex-2011-84)

¹⁴ See EIA22 for a survey.

¹⁵ See http://www.batstrading.co.uk/resources/participant_resources/BATSEuro_Tick_Size_Paper.pdf. London Economics have calculated the economic impact of certain tick size changes through the effect this would have on bid-ask spreads and hence on the cost of capital, EIA22

¹⁶ Trading patterns, liquidity, and the Citigroup split. ITG group: http://www.itg.com/news_events/papers/CitiSplit2.pdf

higher and volume lower; and interestingly, order flow toxicity was lower¹⁷. What is not yet established are the effects on traders or on the issuer.

While the literature carefully delineates the impact of tick size changes on traders and markets, there is no proper analysis on who should make the tick size decision and whether the minimum tick size should be the same for all firms or all trading venues. The current approach in Europe of allowing each venue to choose its own minimum tick size has a variety of merits, but it can lead to excessive competition between venues. The US approach of a uniform minimum tick size removes this problem, but there are deep concerns that it is leading to insufficient liquidity for less frequently traded stocks¹⁸.

3.5 Conclusions

Tick size policy can have a large influence on transaction costs, market depth, and the willingness to provide liquidity. The current approach of allowing each European trading venue to choose its own minimum tick size has merits, but can result in unhealthy competition between venues and a race to the bottom. A uniform policy applied across all European trading venues is unlikely to be optimal, but a coherent overall policy for minimum tick size that applies to subsets of trading venues may be desirable. This coordinated policy could be industry-based such the one agreed to by FESE members.

4. Obligations for market makers

4.1 The measure and its purpose

Obligations for market makers are requirements that a person (or more controversially, a computer program) acting as a market maker must post prices to buy and sell at competitive levels at all times the venue is open and regardless of market conditions. This could be applied to traditional (human) market makers or to algorithmic market makers or both¹⁹. The purpose of this proposal is to improve continuous liquidity provision and to ensure that market makers are actively quoting competitive prices during periods of market stress.

Market makers provide liquidity to traders by being willing to buy when a trader wants to sell and to sell when a trader wants to buy. For providing this service, the market maker earns the bid-ask spread. In the case of large,

¹⁷ Order flow is toxic when it adversely selects market makers, who may be unaware they are providing liquidity at a loss. The fall in order flow toxicity suggests that adverse selection also fell after the reverse stock split.

¹⁸ A further issue, as discussed by Angel (EIA7), is that suboptimal tick sizes may lead to stock splits or reverse stock splits by companies who have a different view of what the tick size should be in relation to the price. This splitting phenomenon seems to be more common in the US, which has a relatively rigid tick size policy.

¹⁹ The current MiFID II draft mentions such a proposal in Article 17 Clause 3.

actively traded issues held by both institutions and retail traders, market makers typically earn sufficient returns to justify the time and capital needed to perform this function. For less actively traded stocks, this may not be the case, and wide spreads, reduced depth and general illiquidity may result. In times of market turmoil, market making in any stock is often unprofitable and the withdrawal of market makers can result in market-wide illiquidity. Consequently, market makers have often received various inducements to supply liquidity such as access to fee rebates, or superior access to the order book, or exemption from short sale requirements, or even direct payments from exchanges or issuers. In return for these inducements, exchanges traditionally required market makers to quote bid and ask prices even in times of stress.

In current markets, much of the market maker function has been taken over by HFT in which a computer algorithm is programmed to buy and sell across markets. Such computerised trading often involves the placement of passive orders (i.e. limit orders) and so, like the traditional market maker, the HFT programme is buying from active traders who want to sell and selling to active traders who want to buy. However, unlike traditional market makers, the HFT programme is not committed to a particular venue, and generally does not have access to superior information, although the computer running the programme may be co-located in the exchange. What is now being considered is whether such high frequency market making should also face obligations with respect to provision of liquidity. Among the various obligations being considered are: maximum spread restrictions; percentage time for quotes to be at the inside spread; minimum quoted size; and minimum quote time.

The policy issue is whether regulators should require market maker obligations for any or all types of market makers. Many exchanges already have some obligations for market makers in terms of quoting so the question is whether this needs to be mandated across all trading venues or extended more broadly to the market making function²⁰.

4.2 Benefits

Market maker obligations can improve market quality and hence raise social welfare. Narrower spreads will induce both informed and uninformed traders to trade more, which in turn increases price efficiency and quickens price discovery. To the extent that obligations improve the depth of the market, traders will find it easier to buy and sell, and transactions costs should be lower. Obligations to set competitive prices could help reduce volatility, and

²⁰ For example, NASDAQ has recently proposed to implement the designation of market makers who are compensated by issuers when committing to a minimum liquidity supply. This new programme is under review. See the document and associated comments at <http://www.sec.gov/rules/sro/nasdaq/2012/34-67411.pdf>

requirements to stay in the market continuously could lead to greater liquidity during periods of market stress.

It should be noted that benefits from such obligations are not guaranteed because of the high costs that they may entail. When companies have contracted for market making services (as, for example, on the Stockholm Stock Exchange), Weaver (EIA8) reports increased market quality in terms of trading volume and liquidity measures for their shares. However, the high costs of paying for contracting have deterred many small companies from doing so.

4.3 Costs and risks

Market maker obligations would be unnecessary if providing liquidity was a profitable and relatively riskless undertaking. The reality, however, is far different, and market makers face a variety of situations in which posting quotes exposes them to the risk of large losses. Moreover, even in the best of circumstances, market making is not cost free, requiring both capital and expensive investments in technology to support operations. To the extent that market making obligations are imposed without corresponding compensation, at least some market makers will exit the market, reducing its liquidity.

How to impose these obligations is problematic. Rules requiring market makers to post narrow bid-offer spreads are often unnecessary for large, active stocks where market making is profitable. However, with less actively traded small stocks, the order flow is more naturally unbalanced and the market maker faces substantial risk acting as the intermediary for the bulk of trading. Market maker obligations on such stocks will impose significant costs on market makers, and too little market making will be provided unless they are compensated. Conversely, if market makers are given too much compensation, trading in small stocks will essentially be subsidised. Such a situation characterised trading on the NYSE (before the regulatory changes in the late 1990s) whereby small stocks generally benefited from rules restricting maximum spreads and mandating continuous prices.

Market making during times of market stress is also an extremely risky proposition as requirements to buy when prices are crashing may lead to bankruptcy for the market maker. An optimal market maker obligation should not force a market maker into bankruptcy, so limits as to what is actually required are both necessary and difficult to even define, let alone enforce. Too stringent obligations will transfer losses from traders to market makers, while setting too lax requirements can result in greater market instability.

Imposing market maker obligations on algorithmic market making trading strategies raises a variety of risks. Many high frequency strategies post bids and offers across correlated contracts. Thus, a high frequency market maker may be buying in one market and selling in another. A requirement to post a continuous bid-offer spread is not consistent with this strategy and, if binding, could force high frequency traders out of the business of liquidity provision. With upwards of 50% of liquidity coming from high frequency

traders, this could be disastrous. A more likely outcome, however, is that any requirement would be evaded by posting one quote at the market and the other off the market or for small size²¹. Moreover, what even constitutes market making in this context is unclear. Many high frequency strategies are actually a form of statistical arbitrage (stat arb), buying where prices are low and selling where prices are high. Using limit orders to implement these strategies is akin to market making, but it also differs in a variety of ways. Forcing market maker obligations on algorithmic trading could reduce these stat arb activities and thereby reduce market efficiency. These issues are discussed in more detail in Cliff (EIA19).

Cliff (EIA19) also discusses the challenges of specifying a requirement that would apply to algorithmic-based market making strategies in a logical or enforceable way. Changes to the original MiFID II 17(3) specification (known as the Ferber amendments) mean that market maker obligations would apply only to HFT systems that operate on maker-taker trading venues and for which more than 50%, or a majority, of the system orders/trades qualify for maker discount/rebates. While this revised application may be more feasible to implement, it may also induce trade to move to venues where this regulation is not binding.

4.4 Evidence

The vast majority of empirical studies on voluntary²² market maker obligations conclude that they improve market quality. These benefits are found in a wide variety of market settings, and across different classes of securities such as equities and options. The benefits are especially felt in illiquid stocks, where generating critical mass in a market is an issue. The fact that virtually every major stock exchange has some form of market maker obligation testifies to their usefulness in enhancing market behaviour²³.

However, empirical research finds that traders under market maker obligations generate these benefits in part because they get compensated. This can be in the form of extra rights such as sole view of the limit order book, ability to short shares 'naked' (without holding the stock) or direct compensation paid either by the trading venue or by the listing company. Moreover, there is a surprising diversity of approaches taken towards how

²¹ See EIA22 for further discussion of monitoring and enforcement issues.

²² By voluntary, we mean that these all refer to situations where there is some quid pro quo for the market making obligation in terms of fee rebates and superior access to the order flow. Compulsory market making obligations (i.e. without compensation) have not been studied to our knowledge.

²³ Of the major markets, only the Tokyo Stock Exchange does not involve some form of specified market making. See EIA 8 for a detailed discussion of these issues.

these market maker obligations are structured. EIA8 notes that “the application of minimum obligations for market makers, as well as the mode of compensation, is uneven across markets on both sides of the Atlantic. Some markets impose minimum obligations on market makers for all listed stocks”.

4.5 Conclusions

The current system of exchanges determining how to structure market maker obligations and pay for them seems to be working well for most markets. We think there is less support for policies that impose market maker obligations for a large class of market participants without a thought-through incentive scheme to support it.

5. Minimum resting times

5.1 The measure and its purpose

Minimum resting times specify a minimum time that a limit order must remain in force. The impetus for imposing a minimum is that markets now feature a large number of fleeting orders that are cancelled very soon after submission. This increases the costs of monitoring the market for all participants, and reduces the predictability of a trade’s execution quality since the quotes displayed may have been cancelled by the time the market order hits the resting orders. The nature of HFT across markets, as well as the widespread usage of hidden orders on exchanges, are responsible for some of this fleeting order behaviour. However, frequent cancelling of quotes may also result from abusive strategies including spoofing, layering, and quote stuffing which can undermine market quality or, at the least, create a bad public perception.

As a result, minimum resting times have been suggested whereby a limit order submitted cannot be cancelled within a given span of time. This measure can take a multitude of forms, such as a uniform 500 microseconds across all assets and securities, or a delay that depends on the security and/or general market conditions. It would also be possible to prescribe different minimum resting times on limit orders to buy or to sell, or that adjust to reflect volatility or other market conditions.

5.2 Benefits

Minimum resting times can increase the likelihood of a viewed quote being available to trade. This has two important benefits. First, it provides the market with a better estimate of the current market price, something which ‘flickering quotes’ caused by excessive order cancellations obfuscates. Secondly, its visible depth at the front of the book should be more aligned with the actual depth. This knowledge of the depth improves the ability of traders to gauge the price impact of potential trades. Quotes left further away from the current best bid or offer are less likely to be affected by the measure since the likelihood of them being executed within a short time is small. Nonetheless, minimum resting times might be expected to make the order book dynamics more transparent to the market.

Minimum resting times may also reduce the excessive level of message traffic currently found in electronic markets. Cancellations and resubmissions are a large portion of these messages, and at peak times they can overwhelm the technological capabilities of markets (as seen for example in the recent Facebook initial public offering (IPO) problems on NASDAQ)²⁴. Some authors also suggest that minimum resting times may reduce the profitability and incidence of spoofing, quote stuffing and other illicit practices. While conceptually possible, there is no clear evidence that such market abuses only involve flickering quotes, and for those that do, surveillance and fines may prove a more efficient deterrent than imposing a minimum resting time.

Minimum resting times may also allay concerns that markets are currently 'unfair' in that high frequency traders are able to dominate trading by operating at speeds unavailable to other traders. This notion of 'slowing down' markets is not generally supported by economic analyses, but it does speak to the challenge of inducing participation if some traders, particularly small retail investors, feel that speed makes markets unfair.

5.3 Costs and risks

Liquidity providers post limit orders available for trade within a period of time in return for an expected gain in the form of the bid-ask spread. Providing limit orders is costly since posting a limit order offers a free option to the market which is exercised at the discretion of the active trader. If an active trader has better or newer information, the limit order poster will be adversely selected, buying when the stock is going down and selling when the stock is going up. As with any option, its value increases with time to maturity and with volatility. Thus, forcing a limit order to be in force longer gives a more valuable option to the active trader, and consequently raises the cost of being a limit order provider. The expected result would be an increase in the bid-offer spread or decreased depth as posting limit orders will be less attractive²⁵.

This reluctance to post limit orders will be particularly acute during times of high volatility when the cost of posting the option is naturally increased. This has the undesirable implication that liquidity provision will be impeded just at the times when markets need it most. It also suggests that there could be a feedback effect if increasing volatility triggers orders, further increasing volatility.

A minimum resting time policy may also change the dynamics of the market by attracting more aggressive high frequency traders whose sole aim is to

²⁴ For discussion on the Facebook IPO, see <http://www.nanex.net/aqck/3099.html>

²⁵ EIA21 describes how a minimum resting time would affect different HFT strategies such as market making, stat arb, 'pinging', (sending and typically immediately cancelling an order to see if hidden liquidity exists) and directional strategies. They suggest a similar impact, although the magnitude of any effect would depend on the length of time that is imposed, market conditions, and the liquidity of the instrument in question.

take advantage of the free options. Depending on the length of compulsory resting, those limit orders close to ‘the touch’ (i.e. the current bid or offer) are likely to become stale (that is, no longer at the efficient price) before they can be cancelled. This can spawn ‘front running’ by automated traders who collect the low hanging fruit from such options. In return, the providers of passive quotes will protect themselves against staleness through yet larger bid-ask spreads, or by simply not posting quotes at all. Using the estimates by Farmer and Skouras, the cost of hitting such stale quotes may be as high as \$1.33 billion per year²⁶.

A final argument pointing to larger spreads concerns the nature of market making in high frequency markets. Modern market makers using HFT have, to some extent, replaced capital and inventory capacity by speed. With minimum resting times raising the risk of losses from limit orders, high frequency traders may reduce their market making activities and possibly be replaced by institutional market makers, such as banks. Reduced competition among market makers and their need to earn a return on their capital may drive up transaction costs for end users. Moreover, to the extent that minimum resting times inhibit arbitrage between markets, which is essentially at the heart of many HFT strategies, the efficiency of price determination may be diminished.

5.4 Evidence

The empirical evidence to date is very limited, since there are very few natural experiments to shed light on the costs and benefits of minimum resting times²⁷. In June 2009, ICAP introduced a minimum quote lifespan (MQL) on its electronic broking services (EBS) platform. These quote requirements set a minimum life of 250ms for their five ‘majors’ (generally currency contracts) and 1,500ms in selected precious metals contracts. In public statements, ICAP credits the absence of a major Flash Crash to MQLs, but of course it is difficult to know what would have happened in their absence. To our knowledge, there has been no empirical analysis of the effects of these MQLs on market or trader behaviour.

Until mid-2011, the Istanbul Stock Exchange (ISE) did not allow the cancellation of limit orders during continuous auction mode unless it was the very last order entered into the system. Depending on the structure and participation on the ISE around the switch, there may be some evidence that can be gathered from that event. Academic research is seeking to identify the effects of this rule change, but the results of any study would be vulnerable to criticism that they could be due to some specific features of the

²⁶ These projected losses borne by market makers assume that market makers do not adjust their limit orders, and that aggressive traders can wait until the end of the resting time before they hit the limit order. Competition between latency traders would probably reduce this time, and therefore the resulting profit, by a non-negligible extent.

²⁷ A commissioned study, (EIA3) examined the effects of minimum resting times inside a simulated market. They did not recommend its adoption.

ISE. Clearly, this market is rather different from the LSE or Deutsche Börse, and it is not clear what can be learned from this experiment.

5.5 Conclusions

The independent academic authors who have submitted studies are unanimously doubtful that minimum resting times would be a step in the right direction, in large part because such requirements favour aggressive traders over passive traders and so are likely to diminish liquidity provision.

6. Order-to-execution ratios

6.1 The measure and its purpose

This measure puts an upper limit on the order to execution ratios (OERs), and as such is part of the larger class of restrictions on order book activity restrictions being considered by policymakers on both sides of the Atlantic. The idea of such restrictions is to encourage traders to cancel fewer orders, and thereby provide a more predictable limit order book. It is hoped that such predictability will improve investor confidence in the market. As cancellations and resubmissions form the bulk of market message traffic, this proposal would also reduce traffic and the consequent need for market participants to provide increasing message capacity in their trading systems.

A number of exchanges have some restrictions on messages or the order/trade ratio (OTR). The LSE Millennium system for example, has message-throttling constraints which limit the total number of messages that can come down a registered user's pipes over a 30 second time frame. It also has a message pricing system which penalises excessive ordering strategies. Sponsoring firms are required to apportion a maximum message rate threshold to prevent sponsored users from entering an overly large number of messages. The limit is set as a maximum number of messages per second per sponsored user and is part of the total limit allowed for the sponsoring firm's allocation. So there are sensible exchange-specific measures already in place that constrain the total message flow and price the externality those messages contribute. The policy question is whether there is value in extra regulation to enforce best practice across exchanges and extend this practice across all trading venues.

6.2 Benefits

Receiving, handling and storing messages is costly for exchanges, brokers and regulators. Whenever an economic good is not priced there is a tendency to use more of it than if the user had to pay its actual costs. If the social cost of messages exceeds its private costs, an externality results; the standard solution is to tax messages. A ratio of orders-to-executions essentially does this, and it can serve to align these private and social costs, thereby reducing the number of economically excessive messages. This, in turn, will reduce the need for exchanges, brokers and other market participants to invest in costly capacity.

With fewer quote cancellations, the order book may be less active and traders may find it easier to ascertain current prices and depths. An order-to-execution ratio also may increase the likelihood of a viewed quote being available to trade, partly because passive order submitters would focus more on those limit orders with a higher probability of execution. An order-to-execution ratio may also help curtail market manipulation strategies such as quote stuffing, spoofing and layering. Quote stuffing is when a trader sends massive numbers of quotes and immediate cancellations, with the intention of slowing down the ability of others to access trading opportunities. Layering refers to entering hidden orders on one side of the book (for example, a sell) and simultaneously submitting visible orders on the other side of the book (e.g. buys). The visible buys orders are intended only to encourage others in the market to believe there is strong price pressure on one side, thereby moving prices up. If this occurs, the hidden sell order executes, and the trader then cancels the visible orders. Similarly, spoofing involves using and immediately cancelling limit orders in an attempt to lure traders to raise their own limits, again for the purpose of trading at an artificially inflated price. These strategies are illegal (Trillium Trading in the US was recently fined by the Financial Industry Regulatory Authority (FINRA) for layering) but they are often hard to detect. By limiting order cancellations, an order-to-execution ratio will reduce the ability to implement these strategies.

6.3 Costs and risks

The nature of HFT and market making in fragmented markets naturally implies order cancellations. Algorithmic trading, for example, seeks to reduce trade execution costs by splitting large orders into smaller pieces and sending orders both spatially and temporally to markets. As orders execute or languish, the execution strategy recalibrates, leading to cancellations and resubmissions. Such a trading approach reduces execution costs for traders and leads to greater efficiency in execution. Many HFT strategies (including HFT market making) involve stat arb across markets whereby movements in a price in one market trigger orders sent to other markets. Again, subsequent price movements in any of the markets will trigger cancellations and resubmissions as part of the process of reducing price discrepancies and enhancing market efficiency.

Many order cancellations are a result of searching for hidden liquidity on limit order books. Exchanges increasingly allow submitted orders to be completely hidden, meaning that the 'best' quotes visible on the book are not actually the best quotes available in the market. To find this liquidity, traders often 'ping' or send small orders inside the spread to see if there is hidden liquidity. Because such orders are typically cancelled, a binding order-to-trade ratio would result in less pinging and, therefore, less information extraction at the touch. As a result, more hidden orders will be posted, leading to a less transparent limit order book. A second effect on the book may arise because orders placed away from the touch (the best bid and ask prices) have the lowest probability of execution. In a constrained world, these orders may not get placed, meaning that depth may be removed from the book away from touch. An added difficulty is that the constraint may be more likely to be binding during times of extreme market activity. Brogaard

(EIA1) argues that this will reduce the willingness of traders to post limit orders during volatile times, thus reducing market liquidity provision when it is most needed.

Finally, there is the calibration exercise of where exactly to set any ratio and to what type of orders or traders it will apply. If the upper limit of the order-to-executions ratio is small, then it will stifle legitimate activities and prevent socially useful trading. For instance, exchange-traded funds (ETFs) and derivatives valuations may become unaligned, leading to inefficient pricing. Because of this, the London Stock Exchange has an order-to-trade ratio of 500/1 for both equities and ETFs/ETPs, with a high usage surcharge of five pence for equities and 1.25 pence for ETFs/ETPs. If instead the upper limit is set high enough not to impinge on legitimate order strategies, it may not have much impact on the market either (a point made by Farmer and Skourous (EIA2)). If the intent is to limit manipulative strategies, a specific charge for messages (and greater surveillance) may be a better solution²⁸.

6.4 Evidence

There have been no published academic studies of OERs, and this greatly limits the ability to gauge the costs and benefits of order activity restrictions in general, and order-to-execution ratios in particular. The commissioned study, EIA18, investigates the effect of the introduction of an OER penalty regime on the Milan Borsa on April 2nd 2012. The authors preliminary findings are that liquidity (spreads and depth) worsened as a result of this policy measure. They also find that the effect is more pronounced in large stocks, although they acknowledge some issues with their methodology.

There are a variety of actual market programs that provide some evidence of OER impact. The ICAP has a monthly fill ratio (MFR) requiring that at least 10% of all quotes submitted into the market must result in an execution. Similarly, LSE's Millennium trading system has message throttling constraints and penalties for excessive ordering strategies. Anecdotal evidence suggests that the LSE message policy was not fully effective in that it gave rise to new patterns of trade in low-priced stocks. The LSE has experimented with changes in pricing effective May 4, 2010 whereby, among other measures, the threshold for the high usage surcharge for FTSE 350 securities increased from an order-to-trade ratio of 100/1 to a ratio of 500/1 (which is still the figure in use as of writing). The frequency of order book updates nearly doubled for a few months as a result before coming down again. Unfortunately, we are not aware of a proper scientific investigation of these effects.

6.5 Conclusions

An order-to-execution ratio is a blunt measure that catches both abusive and beneficial strategies. It may not do too much harm if the upper limit is large enough not to hinder market making and intermediation, but to the extent

²⁸ Oxera outline a number of ways in which an OER could be bypassed or manipulated by HFT, whereby any benefits from the rule may be reduced.

that it is binding on those activities it may be detrimental to both spreads and liquidity. It is unlikely that a uniform order-to-execution ratio across markets would be optimal because it depends upon the type of securities traded and the trader clientele in the market. If a ratio could be structured to target those quotes that are considered socially detrimental directly, then it might be a useful tool for combating market manipulation. The absence of research which investigates costs and benefits, as well as the difficulty of actually setting this measure optimally, suggest caution in adopting this approach for the market.

7. Key interactions

In this working paper the different measures available to policy makers have been discussed individually. A number of key interactions between the measures are considered in the section below.

The presence or absence of circuit breakers affects almost all the other measures except perhaps notification of algorithms. The direction of the effect is harder to determine. Having more stable and orderly markets is likely to improve conditions for many traders, but decreasing the probability of execution for limit orders may adversely affect particular trading strategies.

Likewise, minimum tick sizes affect almost every other measure except perhaps notification of algorithms. The tick size affects the profitability of market making and so will affect market maker obligations. The smaller the tick size, the more onerous are the obligations to post competitive bid-offer quotes because the return from doing so (the spread) is smaller.

Minimum resting times and minimum tick sizes may complement each other on passive orders and possibly conflict on active orders. One of the assumed benefits of minimum resting times is a slowing-down of (passive) activity. Larger tick sizes have this effect as they discourage queue jumping and increase the value of being towards the front of the queue. Larger tick sizes make speed more valuable as it improves the chances to be placed towards the front of the queue, but minimum resting times make this more dangerous for the trader. In that sense, the measures are complementary since minimum resting times blunt to some extent the speed advantage granted by larger minimum tick sizes to faster limit order traders. But minimum resting times also make speed for market orders more valuable as the fastest aggressive order will be first in picking off a now stale passive order. If ticks are larger, this opportunity will be more profitable still, albeit rarer. The interaction is therefore complex and ambiguous.

If minimum resting times and lower minimum tick sizes are introduced for a security, the benefits of either measure may not be fully reaped. The reason is that the existence of a minimum resting time by itself tends to increase bid-offer spreads on one hand, and on the other hand as the true value of the security moves (*ceteris paribus* given the minimum tick size), it is more

likely to render resting quotes stale if tick sizes become smaller. This makes 'picking off' more frequent (but less profitable). It follows that given minimum resting times, a reduction in minimum tick size may not lead to a significant reduction in spreads as passive order submitters need to protect themselves against more frequent sniping.

Order-to-execution ratios and larger minimum tick size both reduce traffic and complement each other. Depending on the non-linearities between quote volume, server speed and stability, quote stuffing may become easier given already large volumes of data, in which case a larger minimum tick size (MTS) (which might be expected to lead to less message volume) makes quote stuffing more difficult. Since the order-to-execution ratio is a blunt instrument which may catch useful trading strategies as well, a higher minimum tick size might allow the order-to-execution ratios to be larger and still accomplish its role in reducing quote stuffing without inhibiting market making too much. If it was found that minimum tick size ought to be reduced for some securities because the spreads are artificially large for liquid stocks, then an order-to-execution ratio may help to allow a smooth transition to lower tick sizes without an explosion of messages.

Market maker obligations and minimum resting times clash in the sense that high frequency traders are required by market maker obligations to post limit orders with tight spreads while minimum resting times mean that other high frequency traders take advantage of those stale quotes. This may mean that high frequency traders snipe each other and that in volatile markets much of the trading would comprise high frequency traders trading with each other.

To conclude, the main lesson of these dependencies is to underscore that whatever rules are implemented, they must be carefully calibrated against other parameters, such as the various tick sizes and the exact circuit-breaking mechanisms in the primary exchanges.

8. Conclusions

We have considered a variety of proposals to deal with the new world of computerised trading in markets. In this working paper, we have summarised the views of studies directed toward understanding the impact of these proposed changes. We further summarise our position below.

The desirability of understanding algorithmic trading strategies and their impact on the market is laudable but achieving this through notification requirements, as, for example, currently envisioned in MiFID II, may not be feasible given the complexity of algorithms and their interactions in the market.

Circuit breakers have a role to play in high frequency markets, and they are found in virtually all major exchanges. Because of the inter-connected nature of markets, however, there may be need for coordination across exchanges, and this provides a mandate for regulatory involvement at least in times of

acute market stress. New types of circuit breakers triggered as problems loom rather than after they have emerged may be particularly effective in dealing with periodic illiquidity.

Tick size policy can have a large influence on transaction costs, market depth, and the willingness to provide liquidity. The current approach of allowing each European²⁹ trading venue to choose its own minimum tick size has merits, but can result in unhealthy competition between venues and a race to the bottom. A uniform policy applied across all European trading venues is unlikely to be optimal, but a coherent overall policy for minimum tick size that applies to subsets of trading venues may be desirable. This coordinated policy could be industry-based such the one agreed to by FESE members. The current system of exchanges determining how to structure market maker obligations and pay for them seems to be working well for most markets. Extending those obligations more broadly across markets and to the market making function more generally is problematic.

The aim of a more stable limit order book is laudable, and minimum resting times seem a possible device to achieve that aim. Many of the independent academic authors have submitted studies which are very favourable to a slowing of the markets. Nevertheless, they are unanimously doubtful that minimum resting times would be a step in the right direction, in large part because such requirements favour aggressive traders over passive traders and so are likely to diminish liquidity provision.

An order-to-execution ratio is a blunt measure that catches both abusive and beneficial strategies. It may not do too much harm if the upper limit is large enough not to hinder market making and intermediation, but to the extent that it is binding on those activities it may be detrimental to both spreads and liquidity. It is unlikely that a uniform order-to-execution ratio across markets would be optimal because it depends upon the type of securities traded and the trader clientele in the market. If a ratio could be structured to target those quotes that are considered socially detrimental directly, then it might be a useful tool for combating market manipulation. The absence of research investigating costs and benefits, as well as the difficulty of actually setting this measure optimally, suggest caution in adopting this approach for the market.

²⁹ This includes countries in the European Union (EU), the European Economic Area (EEA) and Switzerland

© **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.bis.gov.uk
URN: 12/1088