Markets Committee

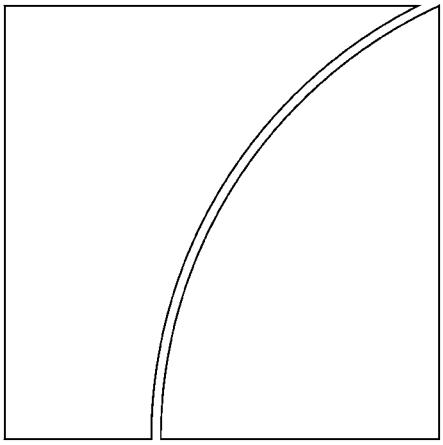

High-frequency trading in the foreign exchange market

Report submitted by a Study Group established by the Markets Committee

This Study Group was chaired by Guy Debelle of the Reserve Bank of Australia

September 2011

BANK FOR INTERNATIONAL SETTLEMENTS

Copies of publications are available from:

Bank for International Settlements
Communications
CH-4002 Basel, Switzerland

E-mail: publications@bis.org

Fax: +41 61 280 9100 and +41 61 280 8100

This publication is available on the BIS website (www.bis.org).

ISBN 92-9131-885-X (print)

ISBN 92-9197-885-X (online)

Preface

In March 2011, the Markets Committee established a Study Group to conduct a fact-finding study on high-frequency trading (HFT) in the foreign exchange (FX) market, with a view to identifying areas that may warrant further investigation by the central banking community. This initiative followed from a number of previous discussions by the Committee about factors contributing to changes in the structure of the global FX market.

The Study Group was chaired by Guy Debelle, Assistant Governor of the Reserve Bank of Australia. The Group drafted an interim report for review by the Committee in May 2011. The finalised report was presented to central bank Governors at the Global Economy Meeting in early September 2011, where it received endorsement for publication.

The subject matter of this report is clearly part of the core expertise of the Markets Committee, which has a long-standing interest in the structure and functioning of the FX market. I hope this report will serve as a timely input to the ongoing discussion about the impact of technological changes, including the rise of algorithmic trading in general and HFT in particular, on the functioning and integrity of financial markets. The FX market focus of this report should also be a valuable complement to a discussion that has so far been based mostly on developments in equity markets.

Hiroshi Nakaso

Chairman, Markets Committee
Assistant Governor, Bank of Japan

Contents

Executive summary

Having come to prominence in equity markets, high-frequency trading (HFT) has increased its presence in the foreign exchange (FX) market in recent years. This development is one aspect of a broader trend facilitated by the wider use of electronic trading in foreign exchange, not only in the broker-dealer market, but also at the customer level. HFT in FX operates on high volume but small order sizes, low margins, low latency (with trade execution times measured in milliseconds) and short risk holding periods (typically well under five seconds). As such, it occurs mainly in the most liquid currencies. While, to date, HFT has been most prevalent among the major currency pairs, it has the potential to spread to other relatively actively traded currencies, including some emerging market currencies.

In equities, where HFT accounts for a significant share of turnover in some markets,[1] the rapid growth of HFT and the perception of predatory practices have generated heightened scrutiny and debate about the benefits and risks posed by this type of trading activity. A number of regulatory initiatives are being considered. A similar discussion is now emerging about the role of HFT in FX.

The assessment of HFT is often hampered by difficulties in identifying this particular type of activity, which is, at times, hard to distinguish from other types of automated (but not high-frequency) trading. There is a lack of reliable data and analysis on the prevalence of HFT as distinct from other forms of automated electronic trading. It is therefore crucial to have a clearer understanding conceptually of what HFT is (and is not) and what it does (and does not do) before assessing the implications of HFT from a policymaker's point of view. Furthermore, given the different nature, structure and size of the FX market compared with equity markets, it is important to ensure that any conclusions about HFT in equities – as well as any regulatory responses – are not inappropriately generalised to HFT in FX.

This report presents the results of a fact-finding exercise conducted by a Study Group consisting of FX market experts from 14 Markets Committee member central banks. Study Group members surveyed existing materials on HFT and also interviewed market contacts (including FX dealing banks, prime brokers, trading platforms and HFT firms) in different financial centres to collect information and views. The objective is to (i) document the facts about HFT in FX and (ii) identify areas that may warrant further investigation by the central banking community.

This report consists of six sections. Sections 1 to 5 constitute the descriptive part of the report. Section 1 describes the characteristics of HFT, how HFT features in the FX market landscape and its relationship with other market participants such as FX prime brokers and major FX dealing banks. Section 2 discusses the effect of HFT on price discovery and liquidity. Section 3 examines the behaviour of HFT in two recent episodes of volatile market conditions. Section 4 highlights the key similarities and differences between HFT in FX and HFT in equities. Section 5 discusses the current practices of self-regulation of HFT in FX.

Section 6 concludes with the lessons learned so far and issues for further consideration:

- *Market functioning*: HFT has had a marked impact on the functioning of the FX market in ways that could be seen as beneficial in normal times. HFT helps to *distribute* liquidity across the decentralised market, improving efficiency, and has narrowed spreads. But the introduction of HFT to the market has affected the ecology of the FX market in ways that are not yet fully understood. Questions remain about HFT participants' willingness to provide liquidity on a sustained basis

[1] The IOSCO (2011) consultation report cites the 2010 estimates by the TABB Group of 56% in the US equity market, 38% in European markets and in the range of 10–30% in Asia-Pacific markets.

under different market conditions. While HFT generates increased activity and narrower spreads in normal times, it may have reduced the resilience of the system as a whole in stressed times by reducing the activity of traditional market participants (eg major market-maker banks) who may have otherwise been an important stabilising presence in volatile environments. That said, recent experience suggests that HFT participants are not necessarily flightier than traditional participants in times of market stress and may be quicker to re-enter the market as it stabilises. Furthermore, the market infrastructure itself, such as the various electronic trading platforms, is also changing in reaction to the growth of HFT and is likely to have a significant impact on how different market participants execute trades over time.

- *Systemic risks*: The 6 May 2010 "flash crash" in equities suggests that systemic risk is perhaps more likely to be triggered by a "rogue" algorithmic trade than by pure HFT, which tends to involve small-size trades, short horizons and diverse strategies. Nonetheless, HFT may under some circumstances accelerate and propagate shocks initiated elsewhere. There are key differences in market structure that may make a flash crash-type event less likely in FX than in equities. But certain longer-term trends such as the adoption of similar technologies have seen FX and equity trading converging.

- *Market integrity and competition*: Many of the "predatory" or "unfair" practices attributed to HFT participants, in the light of their technology-driven ability to detect orders and take advantage of latencies, are in fact not new. HFT is but the latest high-tech, high-speed manifestation of them. A key question is whether other market participants are able to adapt to the presence of HFT, and how the market environment will be affected when those failing to keep up change their trading behaviour or exit the market completely. HFT in FX is subject to three levels of self-regulation. In addition to HFT firms' own risk controls, there is also monitoring by prime brokers. One significant concern here is whether the prime brokers are technologically able enough to keep up with their HFT clients or have the financial incentives to do so appropriately. Furthermore, trading platforms also have rules to help foster an orderly and fair trading environment, but the nature and severity of such rules vary across platforms. At the time of writing, the Foreign Exchange Committees in a number of jurisdictions are also considering implementing enhanced codes of conduct which aim to address the market integrity issues raised by the increased presence of HFT.

- *Looking ahead*: One specific issue for the future of HFT in FX is the potentially differing treatments of electronic trading platforms as a result of the various ongoing regulatory reform initiatives. This is likely to induce some changes to the shape of the FX market. The impact on HFT participants depends on whether there will be more formal regulation of the venues that they currently favour and whether these participants will face some kind of registration requirement.

In sum, HFT in FX is a rapidly evolving phenomenon. It is having a notable effect on the structure and functioning of the FX market, and is prompting behavioural changes in other market participants. All these influence the resilience of the system as a whole, although the impact will continue to change as various participants – including major FX dealers, prime brokers and trading platform operators as well as HFT firms themselves – adapt to the new ecology. Policymakers should continue to keep abreast of this development by maintaining contact and dialogue with the evolving set of relevant market participants.

1. What is high-frequency trading in foreign exchange?

The growth of high-frequency trading (HFT) is one particular aspect of a broader trend in the foreign exchange (FX) market, brought about by advances in information technology and the spread of electronic trading. Before the 1990s, the FX market was predominantly a broker-dealer market. The bulk of transactions took place in the inter-dealer core of the market. Activity between dealers and their customers was in the second or outer tier of this market, where bid-offer spreads tended to be wider than those in the inter-dealer market. Requests for quotes and transactions were typically done over the telephone ("voice").

The advent of electronic broking/trading in the 1990s revolutionised the inter-dealer market. But since this innovation was not yet available in the customer market, the boundary separating it from the inter-dealer market remained (Graph 1, top panel). This boundary blurred when electronic trading became more readily available to FX customers in the early 2000s, when FX dealing banks began to offer trading services to clients via electronic portals (single-bank or proprietary trading platforms) and as the use of credit sponsorship through prime brokerage arrangements grew. Now many types of clients can participate in the over-the-counter (OTC) FX market on a more or less equal footing in terms of price (Graph 1, bottom panel).

Electronic trading can be divided into two main types: (i) *manual*, where instructions are executed by humans on an electronic trading platform; and (ii) *automated*, where instructions are executed by computer algorithms, with little or no human intervention (though still subject to human monitoring).

1.1 HFT is a subset of automated trading

Automated trading, defined as electronic trading using algorithms at some stage in the trade process, has grown rapidly over the past decade and is still evolving. Commonly referred to as algorithmic trading or algo trading, it can be divided into two main strands:

- Algorithmic *execution*: a human trader decides to trade but uses an electronic trading programme to execute the trade. This is often used for larger orders. For example, the programme may use smart order routing to choose where to best trade, or it may use a time- or volume-weighted method to execute the dealer's trade to achieve the best price.[2] Bank traders may use this type of approach to trade via an aggregator; real money investors may use a time-weighted approach to drip-feed a large order to the market.
- Algorithmic trade *decision-making*: a firm builds a model to initiate a trade based on certain key input parameters such as order book imbalance, momentum, correlations (within or across markets), mean reversion, and systematic response to economic data or news headlines. Once a trade decision has been made, the algorithm also executes the trade. Banks' automated risk management tools may also use this method to offset risk automatically. Hedge funds engaged in model-based strategies and specialised HFT funds operate in a similar fashion.

For the purpose of this report, one can think of HFT firms as a subset of *algorithmic decision-makers*. Typically, HFT firms generate earnings from doing a large number of small-size, small-profit trades. The small trade sizes, in part a consequence of operating with low latency (see below), imply that HFT firms take little risk per trade compared with traditional market-

[2] For example, by splitting trades to minimise the footprint on the market.

makers. The risk holding period is also very short, usually well under five seconds and frequently less than one second. As such, HFT requires a liquid underlying market.

Graph 1

Changing structure of the FX market over time

1990s: electronic trading was confined to the inter-dealer market

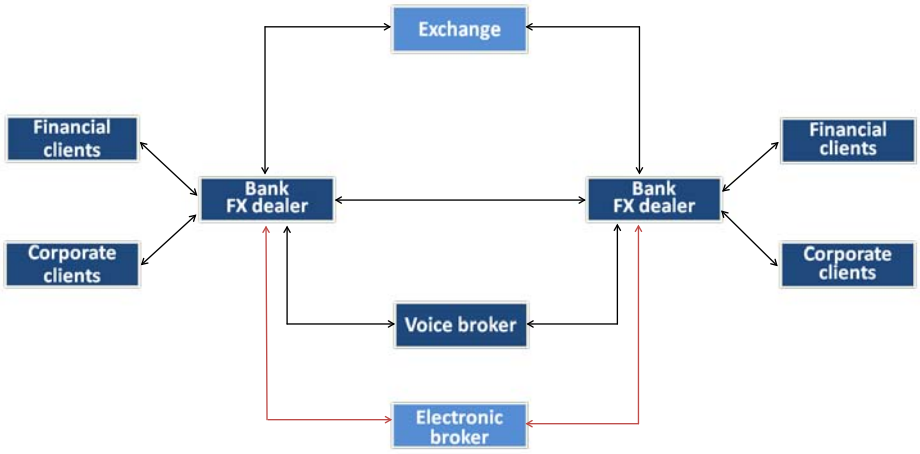

2000s: electronic trading became available to clients; new participants and venues emerged

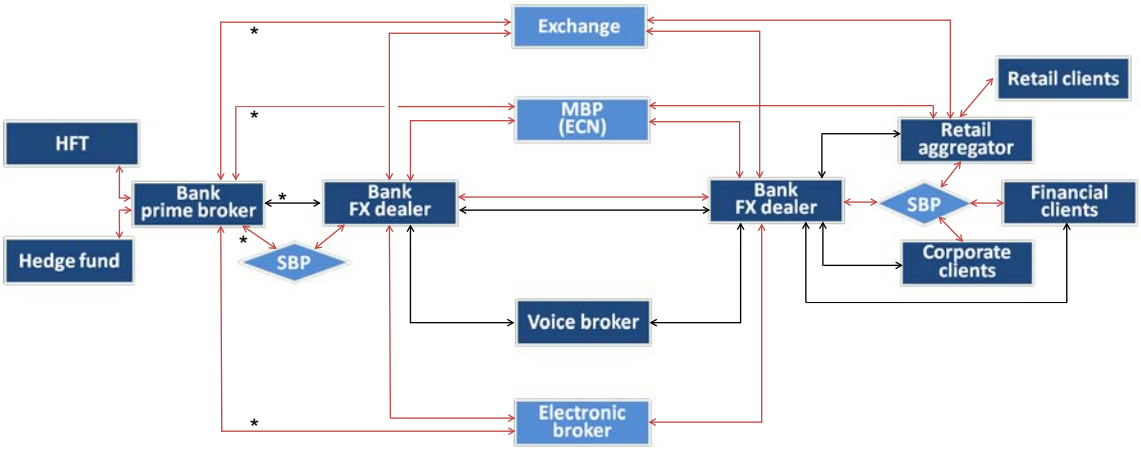

The red lines denote electronic communication; the black lines denote voice communication. HFT = high-frequency trading firm; SBP = single-bank platform; MBP = multi-bank platform; ECN = electronic communications network; Exchange = Chicago Mercantile Exchange, for trades involving FX futures. * indicates prime brokered transactions, which are initiated by the clients but appear (to counterparties) in the prime broker's name.

One of the defining characteristics that set HFT players apart from other algo decision-makers is the high speed with which they detect and act on profitable trading opportunities in the marketplace. Since speed is of the essence, there has been a trend to co-location, ie trading firms moving their servers as close as possible to the trading venue. At the time of writing, market contacts suggest that some HFT participants in FX can operate with latency of less than one millisecond, compared with 10–30 milliseconds for most upper-tier non-HFT participants (for comparison, it is said to take around 150 milliseconds for a human being to blink). Market experts believe that further declines in latency are likely to come at increasingly high costs and may have only minimal financial benefit.

As more and more HFT participants enter the market to exploit these opportunities, or as non-HFT market participants upgrade their systems and reduce their speed disadvantage, the scope to profit purely from the ability to trade with low latency is expected to diminish.

There are indications that this process is already occurring. As a result, some HFT firms are reportedly beginning to rely less only on small-size low-risk trades and to branch out to take on more traditional (directional) risk trades using their sophisticated algorithms (see also below).

1.2 Strategies and business models

A number of different strategies are pursued by HFT firms in the FX market. The unifying characteristic is the method of *implementing* the trading strategies using sophisticated quantitative models and high speed. The various strategies can be classified as follows:[3]

- *Classic arbitrage* exploits the differences between market prices and prices implied by "no arbitrage" conditions. If the price gaps are large enough to cover transaction costs, trades can be executed to lock in a risk-free profit. In spot FX, the arbitrage would be done with a set of currency pairs and the relevant cross rate, eg EUR/USD, USD/JPY and EUR/JPY. This is akin to manual dealers arbitraging USD/JPY, USD/DEM and DEM/JPY in the 1990s on an EBS keypad – but at much higher speed. Arbitrage could also be done across the spot and futures prices of the same currency pair.
- *Latency arbitrage* exploits the small time lag between when market-moving trades take place and when market-makers update the prices they quote. By directly detecting potential price moves, the HFT player can profit from what it has learned ahead of other participants that rely on market-makers' quotes.
- *Liquidity-providing* (or *liquidity-redistributing*) strategies aim to detect order book imbalances for a particular currency pair and pricing discrepancies across trading platforms. The HFT participant earns a spread by arbitraging these differences.
- *Complex event processing* includes a number of different strategies. They aim at detecting profit opportunities by exploiting various properties of currency prices such as momentum, mean-reversion, correlation (with other currency pairs or with other assets) and response to economic data releases.

An individual HFT firm may execute a number of these strategies simultaneously. The quantitative models used may be similar within a particular strategy but can vary significantly across strategies within the one firm. The majority of HFT strategies are designed to benefit from high liquidity and low volatility. Hence there is a tendency for HFT participants to reduce risk when volatility rises. But market contacts suggest that some HFT firms have also developed trading models that are designed to work under more volatile conditions (see Section 3).

Shortly after its emergence in equity markets, HFT appeared in the FX market in the early 2000s. Some equity hedge funds began to apply some of their algorithmic models developed for trading equities to FX, taking advantage of the FX market's very deep liquidity, broad participation, ease of access and arbitrage opportunities. Other firms started out by running pure "latency models", exploiting the different time lags in price updates across different trading venues. Both types had begun to pursue HFT liquidity-providing (or -redistributing) strategies by 2007.

There is a growing trend for HFT firms to contribute prices to trading platforms (market-making, liquidity-providing) rather than just executing on existing prices (liquidity-taking),

[3] Some market contacts have suggested alternative classifications, albeit with similar components. For classification of HFT strategies in equities, see eg Chlistalla (2011) and IOSCO (2011).

although opinion is divided as to how much of this behaviour in FX can be considered "true" market-making (see Section 2).[4] Given their "high-volume, low-margin" business models, HFT firms typically are highly sensitive to the impact of even small errors and exercise tight risk controls (see Section 5).

Market contacts note that as HFT firms become better capitalised, they may extend their risk appetite – for instance, by scaling up already successfully running trading models to gain more volume and increase profits, or by taking on more traditional trade risks.[5] Market contacts also report that some HFT firms are moving to access client flow more directly, either by feeding their price interest directly to banks or to the retail platforms, or even by directly streaming pricing to other market participants.

1.3 Participants and trading venues in FX

HFT participants in FX are mostly specialised independent firms that currently tend to trade only on their own account. Market contacts suggest that several large and better capitalised players account for the bulk of FX HFT volume. There are also a large number of small HFT firms with more limited capital. As noted above, some of these HFT players in FX have evolved from high-frequency trading in equities. Others have been developed by existing FX specialists that have decided to move into the HFT space. A few banks also conduct some HFT in proprietary trading, but they are not major players in this particular space and do not see HFT as an important trend for their business. Rather, they see this as a way to keep up with the technology, which may have positive externalities for their overall FX business.

HFT participants in FX tend to be concentrated in three cities: Chicago, New York and London. Outside these three centres, there are currently very few HFT firms, even in regional FX centres such as Hong Kong, Singapore and Sydney. However, the actual physical location of the HFT firms' offices is irrelevant: what matters is that they co-locate the servers on which they run their algorithms close to the matching engines of the trading venues, which are primarily located in London, New York and Chicago (see below).

HFT firms conduct their FX activities mainly on inter-dealer electronic broking platforms (EBS and Reuters, both London-based companies)[6] and multi-bank electronic communication networks (ECNs, most notably Currenex, Hotspot FX and FXall, typically US-based). They are also active on the Chicago Mercantile Exchange (CME) for trades involving FX futures. The two main inter-dealer electronic platforms were developed earlier (in the 1990s) than the multi-bank ECNs (which were developed in the 2000s).[7]

All venues operate on differing technologies, although there may be more similarities among the newer ones. Participants must adapt to the different technologies, trading rules and trading parameters across venues. The variation in rules reflects, in part, differences in technologies and, in part, different views on market conduct. For example, some venues provide pricing updates at set intervals while others stream prices in real time; the older inter-dealer venues tend to restrict the number of quotes per second and demand certain fill ratios (ie the amount of trades completed relative to quotes submitted), whereas the newer multi-

---

[4] The rebate capture strategy that is widely used by equity HFT firms is a kind of true market-making strategy (see Section 4). But this strategy is so far not prevalent in FX.

[5] Risk (2011) cites market participants as saying that the current tendency for HFT firms to do small-size and low-risk trades with very short, mostly intraday holding periods is not necessarily an inherent property of HFT, but it is the business model that HFT firms have been focusing on so far. There is thus scope for HFT firms to do trades that are akin to those conducted by traditional liquidity providers.

[6] EBS also has matching servers in New York and Tokyo.

[7] See Lee (2010) for more details of the various trading platforms.

bank ECNs tend to allow freer access and appear able to manage higher volumes of data handling. Despite the subtle differences at the front end, much of the architecture is built on common connectivity protocols (APIs) and messaging standards (generally the FIX protocol), as well as straight through processing and Continuous Linked Settlement (CLS).

Market contacts suggest that the larger, more sophisticated HFT players tend to trade on EBS and Reuters, which are currently seen as the predominant source of interbank liquidity in the FX market. However, these wholesale venues traditionally have much larger minimum trade size requirements[8] and tighter trading controls. Smaller players tend to prefer the multi-bank ECNs due to the lower minimum trade size, less tight trading controls (see Section 5) and potentially full anonymity. The fact that some multi-bank ECNs have built-in algorithmic trading functionalities (eg Currenex) also helps to make them attractive venues for very small HFT firms that are just starting up. More developed firms, by contrast, usually use customised in-house models, which provide greater control than do the standard built-in algo functions on ECNs. Market contacts also report that some HFT firms have multiple licences to trade on some platforms. Such multiple presence helps these firms achieve greater market coverage and circumvent certain platform constraints such as limits on the number of quotes that can be submitted per unit of time (see Section 5).

Very few, if any, HFT firms trade solely on single-bank platforms. This is mainly because HFT strategies require a diverse, information-rich (multi-bank, multi-price) environment from which to source trading opportunities. That said, HFT firms do utilise pricing from single-bank providers as one component of their suite of price streams. Since different venues offer somewhat different trading environments (eg due to different trading controls), HFT firms must adapt their trading strategies to the different conditions across venues in order to maintain efficiency (ie they tend to run a portfolio of strategies rather than relying only on one particular strategy).

However, given the greater anonymity in electronic trading (compared with voice, especially at the customer level), it can be difficult to identify what type of player or strategy lies behind a particular trade. Currently in the FX marketplace, there are banks, corporates, funds, institutional investors and even retail users executing trades with some form of algorithm, and some with high frequency. Since HFT firms typically access the various electronic trading platforms through their prime brokers (see Section 1.4), participants on these platforms can usually see, at most, only the prime brokers' names and not their clients' names.

1.4 Relationship with other market participants

1.4.1 FX prime brokerage relationships

The prime brokerage (PB) arrangement is central to algorithmic and high-frequency traders' access to the global FX market. By leveraging the credit provided by the PB along with the accompanying infrastructure support, these clients can gain access to a broad pool of liquidity across various electronic platforms in much the same way as traditional FX market participants can. Accordingly, establishing a PB relationship is among the first important decisions an HFT firm makes after establishing the company's structure. This decision is driven most often by a combination of a prospective PB's cost structure and the suite of services it offers. These factors have important implications for the HFT business model.

[8] EBS also piloted in early 2010 a new service (EBS Smalls) that allows users to trade selected major currency pairs in smaller amounts and increments (100,000 units of base currency, compared with the normal minimum of 1 million units).

There are a handful of very large PB service providers in the FX market. These are typically large investment banks, some of which are also major FX dealing banks. While many PBs indicate that they have a differentiated client base that includes retail investors, hedge funds employing algorithmic execution strategies and high-frequency transaction participants, others say that they have focused more directly on capturing the business of HFT players. Some market contacts suggest that the policies and procedures for taking on board new clients can vary across PB firms.

The PB is responsible for managing the credit it provides to its clients and thus requires the ability to measure and monitor the provision and use of credit through this channel. Some PBs indicate that the nature of the credit relationship is very different depending upon whether the client is an HFT investor or a traditional investor that uses algorithms (not high frequency). As noted earlier, HFT clients typically transact rapidly in smaller trade sizes but tend to hold risk only very briefly. Thus, the outstanding credit utilised tends to be small. By contrast, more traditional participants using algorithmic executions tend to draw on the PB's credit for a longer period of time. From a credit management perspective, the latter client type requires more credit maintenance-related activity due to margin calls, etc. That said, the high-speed nature of HFT means that risk positions can accumulate quickly, raising the need for PBs to become sufficiently speedy themselves in their monitoring of HFT clients (see Section 5).

As noted above, PBs have different strategies for engaging with HFT firms. Their views on the role of HFT firms tend to vary, depending on whether they are courting HFT business. For those committed to supporting the business, HFT firms are viewed as having an important role in the market and are represented as a stable source of liquidity, able to act more nimbly owing to their superior technological capabilities. For those who indicate they are not courting this market segment, concerns are raised about the impact these participants have on the market's liquidity, efficiency and transparency. This variation of views along the lines of business interest is also apparent among other traditional FX market participants (see below).

In addition to the major investment banks providing prime brokerage services, secondary credit sponsorship channels have begun to grow. In this arrangement, known sometimes as "prime of prime", a PB client may be supporting other clients by building on its own PB relationship. This way, smaller firms that may not have qualified as a direct client for the original PB provider may ultimately leverage that PB's credit indirectly, through sponsorship from a client of that same PB.

In the light of the growth of FX prime broking in recent years, a number of market contacts have expressed concern that the prime broking part of an investment bank may be underpricing the risks entailed in providing PB services or the potentially detrimental effect the PB customer may have on other parts of the bank's business. This detrimental effect may take the form of spread compression, which adversely affects the bank's market-making activities. Alternatively, it can take the form of reputational risk since the customer can trade in the PB's name and any inappropriate trading activity undertaken by the customer will be associated with the PB. This concern, however, is not particular to the provision of services to HFT customers alone, but rather a trend across the PB sector more generally.

1.4.2 Relationship with traditional liquidity providers

The growth in electronic trading in general, and the rise of algo trading and HFT in particular have significantly altered the FX market landscape. There is pressure on traditional market participants, notably the major FX dealing banks, to keep up with the technology and speed of their clients and other market participants in order to maintain competitiveness. While some think that this development supports an efficient marketplace, others raise concerns about a market increasingly geared towards algorithmic traders at the expense of manual traders. The March 2011 launch on EBS of decimalised/fractional pip pricing (so-called

"tenths" or "the fifth decimal") for the major currency pairs – a feature that has already existed for some time in certain other single-bank platforms and ECNs[9] – is another recent case in point: while algorithms can relatively easily handle the extra digit, human traders find it more difficult to adapt. The finer pricing also tends to reduce the cost of implementing HFT strategies.

Banks have mixed views about HFT. Having traditionally profited from the spreads they charge for taking risks, FX dealing banks unsurprisingly do not welcome the compression of spreads associated with the rise of HFT (see Section 2), which they see as not fully reflecting the true risk of trading. The high speed and very short holding period of HFT strategies have also drawn criticism, as banks find their own order handling impaired by HFT firms because these firms can detect the order in the early stages of its execution.[10] The introduction of fractional pip pricing further increases the scope for HFT players to out-price banks. Banks need to spend considerable resources to monitor customer relationships closely and screen out any predatory (and thus loss-yielding) practices by HFT firms.

Some banks believe that HFT firms can be valuable as clients. They think it is in their interest to have a diverse set of clients, and they can use the constant flows generated by HFT firms to facilitate and refine their own risk management. Other banks, by contrast, prefer to stick to providing electronic pricing and trading services to their existing clients rather than soliciting business from HFT clients, which they see as opportunistic liquidity-takers. The difference in views probably reflects, at least in part, the differing ability of banks to undertake the investment in systems necessary to manage HFT clients.

Over time, banks have developed various ways to cope with the rise of algo trading and HFT. Banks have made better use of algorithms to execute trades more smartly. They can evade detection by HFT participants in the order book by executing large flows in less transparent venues, including reverting to transacting bilaterally over the telephone.[11] Banks can also internalise more customer flows to capture the spreads for themselves, instead of laying the risk off in the market. Such countermeasures to maintain competitiveness entail considerable spending on information technology, which in turn also affects banks' bottom lines. Those that are less able to make such investments have found their ability to conduct business compromised.

That said, HFT firms ultimately still need the presence and activity of other players such as FX dealing banks, who continue to take larger risk positions, in order to trade. The reverse is not true: dealing banks do not need HFT firms to conduct their business. There may thus be a limit on how much HFT participants can compress spreads or trade at the expense of the banks before the banks would retaliate by providing less favourable quotes to these participants or refusing to trade with them altogether.

---

[9] These do not include Reuters, which is still quoting four digits at the time of writing. In an interview with *FX Week* (29 June 2011), the global head of treasury at Thomson Reuters suggested that decimalisation on this platform was not likely for the foreseeable future.

[10] Commentators often call this phenomenon "front running", evoking the illegal trading practice that involves the use of insider information. However, in the context of HFT, the ability to outrun other market participants results mainly from the use of sophisticated algorithms and high computing speed to rapidly detect potentially market-moving trades, not from the use of insider information.

[11] Published EBS monthly data show that overall electronic FX volume, while having gradually recovered after the financial crisis, remained (as of the first half of 2011) barely comparable with volumes in early 2007 and below the pre-crisis highs recorded in late 2007 and early 2008. While this is not conclusive proof that banks have reduced trading on major platforms to evade HFT players, it is at least consistent with the narrative.

1.5 Estimated scale of activity

Most survey data on FX turnover – whether from global or regional surveys – do not allow direct measurement of HFT activity. Since HFT firms typically gain access to dealers via prime brokers, the dealers participating in these surveys can usually only report the trades with the prime brokers and cannot discern whether the trades contain any HFT-generated flows. Data on prime-brokered business, where collected, are also not broken down to such an extent that makes it possible to identify HFT versus non-HFT activity. These surveys do have some information on execution methods, but cannot yet distinguish automated trading from other forms of electronic trading.

The two major electronic trading platforms (EBS and Reuters) do have data that can distinguish automated trades from manual ones, but cannot directly identify HFT versus non-HFT trades within automated trades. These platforms can in principle construct an estimate of HFT volume by adding up volumes from counterparties that are known to be HFT firms.

Notwithstanding the shortcomings of existing data, some indirect information can still be gleaned from them, given the known characteristics of HFT participants: their trades are necessarily electronically executed; they tend to focus on spot (especially the most liquid pairs); and their activity is geographically concentrated in the United States and the United Kingdom, where the servers of the key trading platforms are based.

The BIS Triennial Survey, for instance, contains some indirect evidence of the rise of HFT as one of the drivers of the rise in FX turnover in recent years. In April 2010, daily average global FX market turnover was $3.98 trillion, up by $657 billion from April 2007. About three quarters of this overall growth came from the rise in spot turnover. Most of this growth was recorded in the United Kingdom and, to a lesser extent, the United States.[12] An estimated breakdown of the growth in spot transactions in these two countries by execution method suggests that trading on multi-bank platforms drove almost half of the 2007–10 spot turnover increase in the United States (trading via electronic brokers contributed to well less than one fifth).[13] The multi-bank ECNs most used by HFT firms are US-based. By contrast, in the United Kingdom the growth in turnover via electronic brokers outstripped that on multi-bank platforms and accounted for over one third of the increase in spot activity there. Reuters and EBS, also well used by HFT players, are based in London.

The BIS Triennial Survey data can also be used to calculate a rough upper-bound estimate of HFT turnover (see Table 1). As a first step, assuming that HFT participants engage predominantly in spot trades, then HFT activity should be captured within the amount of spot trades executed via electronic methods (ie electronic brokers, multi-bank trading systems and, to a lesser extent, single-bank trading systems).

However, this subtotal of electronically executed spot turnover ($872 billion according to the April 2010 BIS Triennial Survey) does not differentiate between automated and manual electronic transactions. Contacts from Reuters suggest that auto-quote trade accounts for 40–65% of total trading, depending on the currency pair. EBS data show that auto-quote-type trading accounted for about 45% of volume on average in 2010. Assuming that this 45%

[12] See King and Rime (2010). Detailed results of the April 2010 Triennial Survey are available at www.bis.org/publ/rpfxf10t.htm.

[13] Information on turnover by execution methods was collected in both the 2007 and 2010 BIS Triennial Surveys. However, the full counterparty details needed to properly adjust for *local* inter-dealer double-counting (which is important for calculating country-level aggregates, or the so-called net-gross basis) were collected only in 2010. Thus, the adjustment for local inter-dealer double-counting for the 2007 data was estimated. Also, owing to different methodologies and coverage, the US turnover growth estimate shown here may differ from the growth figures calculated from the FXC survey for North America.

share applies uniformly to all electronic execution methods, we obtain an estimated $393 billion daily global turnover of automated spot trades for April 2010.

This global proxy is likely to overstate HFT turnover not only because HFT is only a subset of automated trading but also because HFT activity is not global. A similar calculation using data from the North American (FXC) and UK (FXJSC) surveys could yield more geographically plausible proxies. In all cases, the resulting figures turn out to be about 24–30% of spot turnover, in line with market participants' estimates.[14]

Table 1

Estimation of automated spot turnover

Average daily turnover for April 2010, in billions of US dollars

	BIS – global	FXC – North America	FXJSC – United Kingdom
Total, all instruments, all methods[1]	4,124	898	1,900
Of which: spot, all methods	1,590	498	711
Of which: spot, electronic methods	872	264	461
Electronic brokers	*414*	*96*	*268*
Multi-bank platforms	*231*	*168*[2]	*105*
Single-bank platforms	*228*		*88*
Estimated automated spot turnover, assuming 45% share of all electronic	393	119	207
As share of total spot turnover	24.7%	23.8%	29.2%

[1] From the "by execution method and instrument" table of the respective surveys. This total does not correspond to the "headline" total turnover due to issues such as differences in aggregation procedures, incomplete allocation by execution method, and rounding. [2] Includes turnover on single-bank platforms and multi-bank platforms.

Sources: BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity; FXC Semi-annual Foreign Exchange Volume Survey; FX Joint Standing Committee Semi-annual Foreign Exchange Turnover Survey (revised results); BIS calculations.

Two further points of note regarding automated trading or HFT on the two major electronic trading platforms:

- Contact with Reuters indicates that the share of auto-quote trading is also quite high for currencies such as NZD and MXN. This suggests that automated trading, including HFT, has the potential to venture beyond the major currencies.
- EBS suggests that HFT names account for a large majority of the prices made by algorithmic traders on the platform. HFT makes up a much larger proportion of ticket volumes than the nominal turnover volume shows. Some bank contacts believe that

[14] Lee (2010) estimates that HFT in 2010 was 25% of all FX activity (not explicitly defined). King and Rime (2010) report market participants' estimate of 25% of spot FX (source not cited). Discussion with market contacts in Hong Kong suggests that HFT in the narrowest sense is about 10% of total FX activity, or 20–30% if trades with holding periods that last up to a few minutes are included.

HFT makes up a large proportion of EBS turnover; some suggest that HFT firms are among the top customers of EBS.[15]

Finally, some smaller platforms claim to have volumes that can at times rival EBS and Reuters, in part attributable to the high participation of HFT players. For example, market contacts cite Currenex as a platform that does not impose minimum quote life or minimum fill ratios and may therefore be attractive to some HFT firms wishing to trade in a less constrained environment. There is reportedly considerable HFT-to-HFT trading on this platform.

2. Effect of HFT on price discovery and liquidity in FX

2.1 Views of market participants

Notwithstanding complaints about the competition brought about by HFT, bank contacts generally concede that the increased presence of HFT firms has had some beneficial effects on the market: spreads are undoubtedly tighter and some HFT firms appear to enhance liquidity in normal market conditions. Some contacts agree that such firms are efficient in mobilising liquidity around the system and between different venues, which is important in a fragmented marketplace. They note that price behaviour has become more granular, with fewer big-figure moves and instances of gapping than in the days of manual-only trading.[16]

There are, however, questions over the *quality* of liquidity available at a given point in time. Some market contacts argue that HFT provides liquidity at "top of book", for small order sizes, but is detrimental to liquidity for larger order sizes. In particular, since the introduction of the "fifth decimal" for major currency pairs on EBS early in 2011, which reduced the potential cost of providing marginal HFT liquidity,[17] there have been claims that the depth of book has been impaired. One bank contact quantifies this point by noting that spreads in EUR/JPY, for example, have narrowed by 20% for deals that are smaller than €5 million but widened for larger deals (eg two to four times for deals amounting to €20 million).

Indeed, the quality of a bid or ask is determined by more than just its price. The *size* of the bid/ask quote is also an important attribute that determines the ease of transaction, as are the average *"lifetime" of a quote* (ie how long a quote stays in the market before it is cancelled) and the average holding time of positions for market-makers (*"inventory cycle"*) until they close their positions by conducting offsetting trades.

HFT participants typically place a bid or ask only for a very short time, often measured in milliseconds. While HFT players can use their speed advantage to "produce" many quotes, not all quotes necessarily reflect a genuine intention to trade ("quote stuffing", which results in noisy price signals and could slow down price discovery for other market participants).[18] Moreover, HFT-generated quotes can show up simultaneously in multiple trading venues,

---

[15] Risk (2011) also reports dealers as saying that there are four algo traders among the top 10 clients (it is not specified whether these are pure HFT firms, but this seems to be implied given the article's usage of the term "algo"). The article also quotes EBS as saying that there are four algos among the top 15 clients.

[16] This improvement in price behaviour is arguably due not only to the presence of HFT but also to the generally improved ability of all market participants to evaluate data and to trade, given advances in information technology.

[17] For EUR/USD, for instance, the addition of the fifth decimal reduced the potential incremental cost of unwinding a €1 million trade from $100 (1 pip) to $10 (0.1 pips).

[18] However, some trading platforms do impose controls to mitigate this type of behaviour (as well as excessively short quote life) in order to ensure that prices are made only with genuine interest to trade (see Section 5).

giving an impression of good market-wide liquidity, but disappear across venues once one of them is hit.[19] HFT quote sizes also tend to be small, which is not useful for market participants with larger orders to trade. HFT "liquidity providers" also tend to operate with much shorter inventory cycles than do their traditional counterparts. An order executed by an HFT is typically sterilised by a counter-order within seconds. Indeed, the rise of HFT has been associated with a significant reduction in quote life, quote size and inventory cycle, and an increase in the number of transactions. While detrimental to some market participants, this development has clearly been beneficial to those participants trading in smaller amounts of foreign exchange.

That said, a number of market contacts report that they are in fact more concerned about the price and liquidity effects associated with the increased use of algo *execution* methods (rather than algo *decision-making*, including HFT). Some participants in the real money sector have begun to show interest in executing FX orders via bank or third-party algo tools.[20] Algo execution orders tend to be larger in size and unidirectional in nature, and are executed at low frequency. The price and liquidity effects they have could be more persistent than those associated with pure HFT (HFT tends to trade in both directions and carry relatively small risk positions).[21] Nonetheless, both execution and decision-making algorithms have the *potential* to generate herd-like behaviour, to the extent that they may have incorporated similar parameters, which may trigger similar behaviour under some circumstances.

2.2 Findings in the empirical literature

Formal empirical investigations of the effects of HFT on market quality are still scant. The often cited contributions of Chaboud et al (2009) and Hendershott et al (2011) discuss effects of algorithmic trading in general, not HFT in particular.[22] Only recently has a small literature on HFT emerged, using transaction-level data, although this new literature has so far studied only equity markets (see Appendix).

Overall, the empirical literature suggests that algorithmic trading and HFT are neutral to beneficial for market quality, in that volatility has declined and quoted spreads have narrowed. The findings on liquidity provision are mixed. Generally, algorithmic liquidity appears to be more strategic than non-algorithmic liquidity and tends to decline when volatility rises. Regarding price discovery, contrary conclusions have been drawn for FX and for equities. Taken as a whole, the literature consistently shows that information in orders has become more relevant than information in trades.

The literature highlights the importance of accounting for endogeneity: not only does HFT activity influence market conditions, it also responds to changes in market conditions.

---

[19] This is often referred to as a "liquidity mirage". Again, this behaviour need not be attributable only to HFT but can be prevalent among other market participants too.

[20] For example, AES, an algorithmic execution tool that operates as an agency business for Credit Suisse, has been growing strongly in FX since 2008. Nominal turnover has grown around fivefold since 2008. Of the clients that use this service, about 60% are global macro investors, 20% are asset managers and 15% are proprietary trading desks, with the remainder (5%) including corporates, pension funds and sovereigns. Other banks are beginning to offer similar tools. Investment spending in this area is said to be strong at the moment. Some multi-bank ECNs are also said to provide similar algo execution tools, or are building them.

[21] See CFTC-SEC (2010) for a description of how this point on the persistent effects of algo execution played out in the US equity market during the 6 May 2010 flash crash.

[22] The literature usually defines algorithmic trading as the use of computer algorithms to automatically make trading decisions and submit orders. It thus includes HFT as well as order flow originating from aggregation, risk management, agency execution algorithms or any proprietary directional trading models.

Critically, the literature suggests that the effect of HFT depends on the strategies chosen. At the time of writing, published empirical results exist only for strategies in equities. Given the differences in market structures (see Section 4), HFT strategies in FX may well differ from those in equities. Accordingly, results derived from the study of HFT in equities may not necessarily apply in FX.

Chaboud et al (2009) is the only major contribution that investigates algorithmic trading in *FX markets*. Using EBS data ending in 2007, the authors provide four stylised findings:

- Algorithmic ("computer") orders are more highly correlated than human orders: the probability of computers trading with humans is higher than in a pure random setting, whereas the probability of computers trading with computers is lower.
- There is no evidence that algorithmic trading causes an increase in FX volatility in the sample period. There is some evidence that it may have caused a slight reduction in realised volatility.
- Orders by human traders have a more important price impact in EUR/USD and USD/JPY than algorithmic orders. But the impact is about equal in EUR/JPY, consistent with the notion that triangular arbitrage by computer traders plays an important role in the cross rate.
- During the minute following the US non-farm payrolls announcement, when volatility is high, the share of liquidity provided by algorithmic traders decreases, relative to the share of liquidity provided by human traders. During the following hour, however, algorithmic traders increase their liquidity more than human traders do. In providing liquidity, computer traders are found to react more strongly to changes in order flow than human traders. This suggests that computer traders provide liquidity primarily when an adverse shift in prices due to new information is unlikely. Computer traders are thus said to behave more strategically than human traders, in the sense that they are more successful in avoiding trading with informed counterparties.

The authors conclude that the growth of algorithmic trading has not lowered market quality and emphasise that ultimately the trading strategy will be relevant for market quality, not the frequency at which orders are placed.

Empirical studies on the effect of algorithmic trading and HFT on the equity market are discussed in the Appendix.

3. Behaviour of HFT in FX in times of stress

Available data suggest that algorithmic investors, including HFT, have added significant volume to the FX marketplace in recent years. Still an open question, however, is how the liquidity associated with this added volume differs from that provided by traditional market participants (see Section 2). There are questions about the stability of HFT players as liquidity providers in stressed market conditions, particularly since these participants are highly sensitive to risk and tend to hold positions on an intra-day basis only. As bid-ask spreads widen out and prices begin to gap, some observers suggest that HFT participants may be more inclined to pull out of the market, at least temporarily, especially once their net positions are close to neutral. HFT participants are not under any obligation to make markets, and they tend to seek to exploit only very small moves and are thus reluctant to open themselves up to the risk of large adverse price movements.

However, others argue that the superior technology of HFT firms enables them to nimbly navigate such difficult market conditions, making them less likely to exit in the face of volatility. According to market contacts, the specific strategies that HFT firms deploy can be a function of market conditions. In volatile markets, some firms may switch from trading models

that are more suited for normal market conditions to those that have been deemed effective in periods of volatility.

A closer look at two recent episodes of high market volatility helps to shed light on the behaviour of HFT players. One key question is whether HFT-provided liquidity withdrew at all and, if so, whether it withdrew more or less than did liquidity from the more traditional sources.

In discussing these events, it is important to keep in mind that they occurred in an environment of heightened volatility across many markets, not just in FX. Indeed, liquidity in the foreign exchange market, while adversely affected, was generally still superior to that in many other markets.

3.1 FX market on the day of the "flash crash"

During the early afternoon of the New York trading day on 6 May 2010, sharp price action was witnessed in the equity markets, as well as in other markets, including the FX market.[23] EBS data on 6 May show that algorithmic execution comprised about 53.5% of total activity (versus 46.5% manual) on that day, which was higher than on average (45% algorithmic, 55% manual for 2010). The higher than normal share of turnover on 6 May suggests that algorithmic participants did not reduce activity more than did traditional market participants. That said, there are no refined data on the behaviour of HFT accounts as a fraction of algorithmic activity. Anecdotally, EBS contacts suggested that HFT did remain active throughout the volatile session. However, some market contacts, also citing anecdotes, suggested that HFT activity tends to gravitate towards the major liquidity-providing platforms in volatile market conditions. Thus, EBS and Reuters could be among the last venues from which algo traders, including HFT participants, would exit.

Conversations with dealers reveal a mixed picture about the stability of algorithmic players' engagement with the market on 6 May. Some FX prime brokerage providers have suggested that HFT players, in particular, did remain active in the marketplace, with some citing their ability to effectively weather such volatility given their technologically superior risk management capabilities. Others suggested that while there may have been some decline in HFT activity, this was consistent with the behaviour of traditional liquidity providers such as banks, who were also seen to be reducing activity, becoming more risk-averse and widening spreads significantly as developments unfolded. Yet others have suggested that algorithmic accounts in fact did pull back amid the increased volatility.

Raw data from EBS provide a more nuanced view of activity in the EUR/USD pair – including some indirect indications of HFT – during this particular episode of sharp price action on 6 May. Graph 2 plots the price action in the EUR/USD pair on that day between 14:30 and 15:30 New York time against two additional series: (i) the ratio of algorithmic investor order submissions on 6 May to average algorithmic investor order submissions for the prior period; and (ii) the ratio of manual investor order submissions on 6 May to average manual investor order submissions for the prior period.[24]

---

[23] For reference, in equities, HFT participants trading the S&P 500 futures contract were initially blamed for the sharp falls in the prices of many equity-linked instruments and the underlying shares. However, the subsequent official investigation concluded that HFT participants did not cause the crash, did not seem to have altered their strategies on that day and, most notably, did not exit en masse even as market volatility rose (Kirilenko et al (2010) and CFTC-SEC (2010)). Evidence from FX presented in this section also suggests that algorithmic and HFT players did not completely desert the market.

[24] The prior period is defined as the same time of day on the previous 25 trading days. The order figures reflect the total number of orders submitted within 10 pips from the best price for each one-minute increment.

Graph 2

Algorithmic and manual quote submissions on 6 May relative to prior period averages

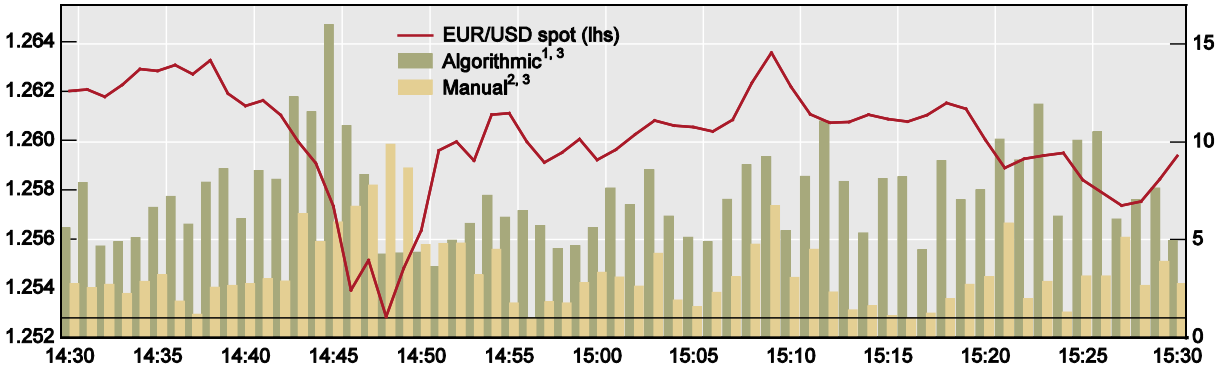

[1] Number of EUR/USD quotes submitted by algorithmic traders on 6 May divided by the average number of EUR/USD quotes submitted by them during the same time of day on the prior 25 trading days. [2] Number of EUR/USD quotes submitted by manual traders on 6 May divided by the average number of EUR/USD quotes submitted by them during the same time of day on the prior 25 trading days. [3] A ratio of greater than 1 indicates higher than average activity.

Source: EBS.

The data show that the number of both manually and algorithmically submitted orders was in fact sharply higher than the average of the prior period. Additionally, the data show that the scale of this increase in the number of orders submitted on 6 May, relative to the average, was more significant[25] for the algorithmic participants than for the manual participants during this particular point in time. Although this measure does not provide insight into the liquidity actually provided by these players, the increase in order submissions is certainly consistent with the higher than average percentage of activity represented by algorithmic investors overall on 6 May, as cited earlier.

Graph 3

Share of PTC-generated algorithmic activity on 6 May relative to prior period average

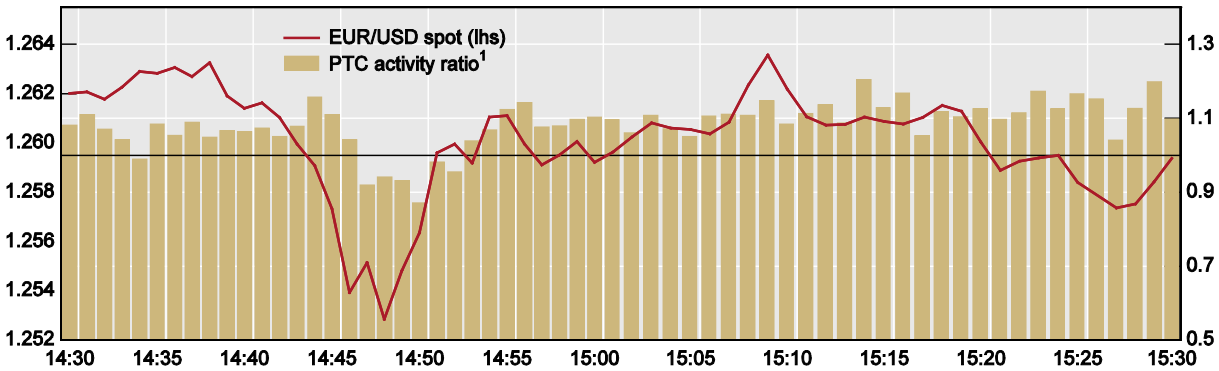

[1] Share of total algorithmic quotes submitted by professional trading community (PTC) participants (ie non-banks) on 6 May divided by the average share of total algorithmic quotes submitted by the PTC during the same time of day on the prior 25 trading days. A ratio of greater than 1 indicates higher than average participation of the PTC.

Source: EBS.

While Graph 2 illustrates higher than average order activity among algorithmic investors on EBS, it does not provide a more discrete look at HFT activity itself, which is one subset of the

[25] However, there is a brief period when this trend briefly reverses, which coincides with the sharp change in trend in EUR/USD shortly after 14:45.

algorithmic category. Graph 3 provides a closer, though still indirect, look by plotting the share of algorithmic activity generated by non-bank participants – or the professional trading community (PTC) – during the same one-hour timeframe on 6 May. These data indicate that PTC activity as a share of total algorithmic activity during this period of volatility was higher than the average over the prior period. This suggests that the increased contribution of activity by algorithmic participants shown in Graph 2 was driven largely by the increased activity of PTC participants.

Of course, PTC activity can include more than HFT and thus these data cannot directly confirm the activity generated by HFT players. Nonetheless, they are at least consistent with anecdotal suggestions that HFT players remained active throughout the session (non-bank HFT firms accessing the market through prime brokers are reported as PTC). In addition, given the larger than average number of orders submitted per minute, as evidenced in Graph 2, the data suggest that the strategies employed by those algorithmic participants who remained active during this period of volatility were more, rather than less, consistent with HFT (eg more frequent submissions of orders per unit of time), although this increase in activity could be attributable to other factors as well, including the generally more volatile market conditions at the time.

That said, HFT firms' continued presence and high trading volume may not be conclusive indicators of the quality of liquidity in times of stress. Experience in the equity market on that day provides an anecdotal illustration. The official investigation of the flash crash finds that *equity* HFT participants initially absorbed some of the selling pressure from fundamental (non-HFT) sellers, acquiring a net long position. But given their short inventory cycles and limited risk appetite, the HFT participants soon had to start selling their long position to reduce market risk, thereby competing for liquidity with fundamental sellers. The CFTC-SEC (2010) report also points out that the rise in turnover during the flash crash was attributable to "hot potato" trading among HFT firms in the absence of fundamental buyers. This increased trading volume in turn amplified the selling pressure from execution algorithms with programmed execution rates linked to market turnover. This anecdotal observation raises concerns that increased HFT activity in volatile markets could in some cases accelerate and magnify price action, resulting in adverse feedback with possible systemic implications (see Section 6.2). Furthermore, the extraordinarily high quote traffic overloaded the NYSE IT systems; quotes were disseminated with delays of up to 20 seconds.

3.2 JPY move on 17 March 2011

Early in the morning of 17 March, USD/JPY declined by some 300 pips from around ¥79.50 to below ¥76.50 in just 25 minutes between 05:55 and 06:20 Tokyo time (16:55–17:20 New York time on 16 March). This move was in part linked to the compulsory stop-loss trades for retail FX margin traders executed by the retail aggregators.[26] The background and timing of the episode are notable in three ways:

- It occurred against a backdrop of historic-high levels of net long USD/JPY positions among Japanese retail FX margin traders, who tend to invest with a contrarian view (USD had been depreciating against JPY for a while).
- It occurred during a less liquid part of the global trading day, at the end of the US session and before the Asia session was fully open.

[26] As from February 2010, the retail aggregators are obliged by regulation to exercise a "compulsory loss-cut rule" such that their clients must cut their positions whenever losses exceed the margin posted. In addition, a regulatory cap has been set for the maximum leverage investors can take against the margin (50 times as from August 2010, lowered again to 25 times as from August 2011).

- It coincided with the time of the day when many OTC retail aggregators have their scheduled system halts to run the batching process to prepare for the next day and to calculate clients' margins for regulatory purposes. Four of the five largest retail aggregators in Japan have system pauses between 05:45 and 06:10 Tokyo time.

Because of the system halts, clients were not able to post additional margins or close positions to avoid triggering automatic stop-outs. When the retail aggregators proceeded with the compulsory stop-outs, it resulted in a wave of USD selling in a relatively thin market. As a result, USD/JPY started to exhibit a free fall. Many banks withdrew from market-making. Those that continued to make markets widened their spreads so much that their bids were far below the last prevailing market price, making it difficult for counterparties to accept and transact. A vicious cycle of USD/JPY fall and stop-losses ensued, until the pair hit ¥76.25 at around 06:20. In the next 30 minutes or so, USD/JPY recovered to ¥78.23 as hedge funds and new retail investors began to build up fresh USD/JPY long positions. Banks, having withdrawn from making prices during the most volatile period, also resumed market-making.

USD/JPY dipped again at around 07:00 to reach ¥77.10. This move coincided with another round of automated stop-outs, executed this time by FX margin-trading brokers that participate on a particular trading platform on the Tokyo Futures Exchange. This platform had a scheduled daily halt between 05:55 and 06:55. When the system restarted at 06:55, numerous compulsory stop-out orders were generated over five minutes to the six market-makers that have obligations to provide prices to this platform (market contacts estimated $2 billion of USD/JPY selling).

Although this episode largely reflects factors that are specific to retail margin trading and the structure of associated systems, it offers another instance of HFT behaviour under volatile FX market conditions. During this episode, both HFT players and traditional market-makers reportedly withdrew, to a large extent, from the market. But in the light of the specific time of day, the withdrawal could also be related to system pauses, rather than just the operation of the HFT firms' internal risk controls. The episode also suggests that, even in trading venues with designated market-makers (as is the case with this particular platform in Tokyo), there is no guarantee of the *quality* of the quotes. Some market-makers with formal obligations to quote prices widened their bid-offer considerably during that time.

4. HFT in FX versus equities

Since the information technology and trading techniques underlying HFT in FX have their origins in HFT for equities, it is not surprising that many similarities exist between HFT in the two markets. For example:

- A focus on the most liquid segment of the electronic marketplace: in equities, primarily blue-chip stocks, index futures and exchange-traded funds; in FX, mainly the most liquid currency pairs such as EUR/USD, USD/JPY, GBP/USD and potentially also some other well traded currencies such as AUD, CAD and MXN.
- Similar impact on markets: narrower spreads, smaller quote size, shorter quote life, shorter holding time, more frequent trades.
- Some potential common concerns: liquidity mirages, potential for "fair weather" market-making, possible artificial "hot potato" volume, incentives for traditional market-makers to go into less transparent venues (see Section 6).

In fact, it could be argued that some aspects in both markets are exhibiting similar longer-term trends or convergence. For instance, both markets are reportedly becoming more order-driven (ie prices determined mainly by the flow and continuous matching of buy and sell orders submitted by participants) rather than quote-driven (ie prices determined mainly by

the bid/ask quotes submitted by market-makers and other participants who are ready to transact at those quotes) (see further discussion in Section 6.2).

However, there are notable differences as well. Given the different market structures, HFT strategies in FX (mainly OTC, global) may differ considerably from those in equities (exchange-based). For example, *rebate capture*, a widely used HFT strategy in equities, is made possible by local exchanges offering rebates to attract order flow and boost volume.[27] An HFT player can profit from this by generating many passive bids or offers for other market participants to trade on. The HFT captures the rebate fee each time an order results in a passive trade, regardless of whether the trade itself makes profits or not. But if the HFT player trades actively to reduce a position by either hitting a bid or lifting an offer, it does not obtain the rebate. This type of pure market-making strategy is so far not prevalent in FX.

Table 2

A stylised comparison of foreign exchange and equity markets

	Foreign exchange	Equities
Market structure	Decentralised, OTC (except for CME)	Organised around formal exchanges
	No formal market-making requirements	With more formal market-making requirements
Trading hours	Essentially 24 hours each day; about five and a half days a week	Fixed trading hours each day; five days a week
Fee/commission	Built into the bid-ask spread	Exchanges can charge fees, offer rebates, etc
Regulation of trading venues	Self-regulated (except for CME)	Formally regulated
Latency	More variable across venues	Relatively predictable
Short-selling	No restriction for freely traded currencies	May be restricted/banned under some conditions
What is traded	Cash for cash (ie relative price)	Cash for equities (absolute price)
Rationale for trading	For both investment and transaction purposes (eg due to trade in goods and services, or foreign assets)	Mainly for investment purposes

The scope for *latency arbitrage*, which exists in FX, may also differ somewhat in the two markets. Latency arbitrage in equities was made more prevalent by the 2005 SEC Regulation NMS (National Market System), which requires brokers to execute each customer trade at the national best bid or offer price (NBBO) across all US equity exchanges at that given moment. The time lag between the appearance of bids/offers on the raw price feeds from the various exchanges (to which HFT firms have access) and the posting of the NBBO on consolidated data feeds (to which most other participants have access) provides a potential arbitrage opportunity for HFT firms.[28] According to one market contact, the transparency of regulated exchanges means that latencies tend to be more predictable. In

[27] Market contacts suggest that, in Canada, the majority of equity HFT activity is from rebate capture.

[28] This time lag essentially allows HFT participants, given their high-speed computing power, to project what the NBBO would be some milliseconds later. HFT participants can then quickly take positions to profit from what they have learned ahead of other slower market participants.

the decentralised FX market, by contrast, transparency is relatively lower (and there is no regulation-enforced equivalent of NBBO); latencies are therefore more variable across platforms. But as these latencies decline and converge across markets over time (eg with improved information technology), FX and equity trading may begin to look increasingly similar.

Another major difference is the presence of "real" demand in and the greater diversity of the FX market. FX is traded not only as an asset in itself, but also due to underlying global trade and capital flows. The FX market includes a broad community of non-speculative real end users with diverse objectives: investment actions of "real money" managers, foreign reserve management actions of central banks and national sovereign wealth funds, hedging actions by global corporates, and retail users. These participants have a variety of holding periods. The FX market is also an essential secondary market underlying primary investment markets such as equities, commodities and fixed income. This breadth and coverage should help to reduce the likelihood of flash crash-type extreme moves in FX markets, especially for the major currency pairs during the more liquid part of the global trading day.

Finally, according to market contacts, while the presence and activity of HFT players have been known to drive traditional market-makers into less transparent venues, trading in "dark" venues has so far been much less common in FX than in equities. There have been discussions among the largest FX banks of plans to set up a more "exclusive" trading venue with rules of engagement that better cater to traditional FX market participants with a genuine intention to deal, although this alternative venue is not supposed to be "dark".[29]

5. Self-regulation: current practices

Since the FX market is largely a decentralised OTC market, not subject to formal regulations, trading practices, including those of HFT firms, are governed mainly by self-regulation.[30] Non-bank HFT firms in FX are subject to three levels of trade control:[31] (i) their internal risk management systems; (ii) the risk management controls of their prime brokers; and (iii) any trade controls imposed by the trading platform. Moreover, to the extent that any specific predatory HFT trading behaviour can be identified, counterparties will seek to limit their trading exposure to those firms generating the predatory trades. This ultimate sanction, of cutting off trading with a counterparty, is possible as long as trading is not completely anonymous (eg the counterparty is identifiable by at least its code on a trading platform, if not its name). That said, the Foreign Exchange Committees in a number of jurisdictions do issue guidance on different aspects of FX trading and related business. At the time of writing, such guidance is being reviewed to reflect recent trends and to enhance the codes of conduct for the FX market, taking account, as needed, of the changes induced by the increased presence of HFT.[32]

---

[29] Apparently in response to this potential move by the major FX banks, EBS piloted in late 2010 and launched in August 2011 a new functionality (available to banks only) called "Continuous Match", which provides manual traders with an easier way to execute large orders. These block trades would take place outside the normal order book so as to limit market impact.

[30] For reference, Clark (2010) discusses internal risk management and pre- and post-trade risk controls for the case of HFT in equities.

[31] Trade control incorporates the management of credit risk, as this impacts the ability of the HFT to trade.

[32] For example, the FXJSC in London is reviewing the 2009 version of its Non-Investment Products (NIPs) Code, the 2011 update of which will include a new section on electronic trading (see the minutes of the Committee's January 2011 meeting). The New York Foreign Exchange Committee is reviewing its 2006 paper on autodealing with a view to updating the document in the coming months.

5.1 HFT firms' internal risk control

As already noted, HFT strategies generate profit by making a large number of small-size low-margin transactions over a short period of time. One cornerstone of their business is the effective reduction or management of their market risk – including the cessation of trading in volatile markets if certain parameters are violated. This is especially important since in most cases the owners' capital is directly at risk. HFT firms' risk management systems are both fast and sophisticated, and potentially superior to those of their bank counterparts. Furthermore, these firms' trading algorithms are extensively backtested prior to actual live trading. According to anecdotal reports, these risk controls were much faster in their reaction to the flash crash than the exchange circuit breakers. Some HFT firms diversify their trading risk by operating a broad portfolio of different trading models or strategies, with each model or strategy having only a small risk allocation (see Sections 1.2 and 3).

5.2 Monitoring by prime brokers

HFT firms access the FX market primarily through a prime brokerage relationship using the credit of the prime broker (see Section 1.4). As a result, PBs have a strong incentive to monitor and control the HFT clients' market access since it is their name, reputation and risk that are directly observed in the market. PBs monitor and control the HFT clients' credit usage and trading performance through their internal risk management systems. This may require co-locating the PBs' risk management servers with the main trading hubs in order to minimise latency. For the risk management process to be fully effective and to minimise the credit risk to the PB, including the impact of a potential rogue trading algorithm, the PBs' risk system must be both real-time and as sophisticated as that of the HFT. Given the short trading time frames involved, risk positions can accumulate rapidly.

Indeed, many prime broker contacts have indicated that they are actively working together with the trading platforms to create a more efficient and timely process for limiting or cutting off credit to clients across the platforms if needed, citing this as a major priority. One such recent industry initiative was launched in late June 2011.[33] It is designed to help PBs monitor their clients' credit risk across different electronic trading platforms by aggregating positions across platforms in real time and providing the capacity to terminate specific clients' access to the market ("kill switch") if unacceptable trading positions are being established. Such real-time capacity to monitor and manage clients' credit could shorten the time lag between unacceptable positions being accumulated and trading being shut down, thereby reducing the overall exposure. While the time lag is reduced, it is not eliminated: Haldane (2011) highlights as an issue the fact that trading and risk monitoring have now become so fast and complex that they have moved well beyond the normal capacity of direct human comprehension.

It remains to be seen whether this and any other risk mitigation initiatives would help address the concern that PBs seeking to generate more income from greater trading volume may be underpricing the services they provide to clients, including HFT clients, relative to the risks they are incurring from that business (see Section 1.4).

[33] Jointly developed by Traiana, a company specialising in electronic post-trade processing for OTC FX and other transactions, and a number of major FX prime brokers and trading platforms. See press release: http://www.traiana.com/news/pressreleases/fullstory/?id=70.

5.3 Trading platform rules and controls

The trading platforms used by HFT firms also have controls to influence participants' trading behaviour. The two main electronic broking platforms (EBS and Reuters) have extensive trade controls, including requirements on the minimum amount of time that a quote has to remain active (minimum quote life, MQL), on the percentage of actual trades conducted relative to the total amount of quotes submitted (minimum fill ratio), and limits on the maximum number of quotes that can be submitted in a specified time interval. These controls vary by currency pair, with the major currencies having a shorter minimum quote life and a lower required fill level.

In addition, both electronic broking platforms have a code or rules of conduct by which participants must abide. The electronic brokers monitor the behaviour of their trading counterparties and adjust the trading parameters as required in order to maintain an orderly market between the traditional non-algorithmic (manual) participants and the algorithmic (including HFT) accounts. For example, following the introduction of decimalisation for major currency pairs in March 2011, EBS announced that it would increase monitoring through the use of a "top of book quote interrupt" measure, which tracks the length of time for which a quote that stays undealt remains at top of book. Ultimately, the platforms have the option of reducing or removing the trading privilege of a predatory HFT counterparty.

Although the electronic broking systems themselves do not have full risk management capability, they allow the PB to specify the gross daily trading volume that each account can trade. However, this does not determine the amount of risk being undertaken by a specific counterparty. The ultimate risk control for the HFT participants' trading on the system lies with the PB.

The CME, an exchange that is formally regulated by the Commodity Futures Trading Commission (CFTC), has controls that are somewhat similar to those of EBS and Reuters. Other ECNs, including the multi-bank platforms, have less extensive trading controls. For example, Currenex does not have requirements such as minimum quote life or minimum fill ratios. The controls that these other platforms do have are there mainly to facilitate more orderly trading when the system is overloaded and thus are typically triggered only when volumes increase sharply.

In the case of single-bank platforms, the bank that sponsors the platform controls the access of HFT clients to the system and can adjust the terms of access for each client based on the performance of each trading relationship. If a particular HFT client's trading relationship is shown to be unprofitable, the sponsoring bank can change the terms (eg size of the bid-ask spread) offered to that client or ultimately cut off access completely.

6. Lessons and issues

The increased use of high-frequency trading strategies in foreign exchange in recent years is a direct result of advances in information technology, the rapid shift from direct bilateral transactions to automated trade execution over the past decade, the proliferation of alternative trading venues and the growth of prime brokerage credit sponsorship. It has nevertheless raised issues for market participants on both the sell and buy sides of the market, and for the global regulatory community.[34] The issues can be grouped into concerns

[34] See Haldane (2011) for one discussion of these issues. IOSCO (2011) discusses market efficiency and integrity issues raised by technological changes, including the rise of HFT; it also mentions a number of initiatives in this area (mostly pertaining to equities) by the regulatory authorities in a number of jurisdictions.

about the effect on market functioning, systemic risk, and market integrity. It is important when evaluating these developments to acknowledge the financial incentives for other market participants, including PB service providers and trading venue operators, to mitigate potential negative outcomes from HFT.

6.1 Market functioning

HFT has had a marked impact on the functioning of the FX market in ways that could be seen as beneficial in normal times, but also in ways that may be harmful to market functioning, particularly in times of market stress. HFT has arguably helped to increase market efficiency by mobilising the power of new technology in the marketplace and by creating incentives for other participants to upgrade their own technology as well. However, the competition posed by HFT firms through their ability to compress spreads and make profits at the expense of traditional market-makers, has resulted in some market-makers pulling back or changing how and where they provide liquidity. That is, the introduction of HFT to the market has affected the ecology of the FX market. This has a number of consequences:

- In normal times, the presence of HFT is beneficial to market functioning in providing liquidity and compressing spreads, at least for smaller trade sizes. But liquidity for larger trade sizes may in some cases have become inferior. Moreover, given its sensitivity to risk, HFT liquidity may evaporate rapidly in stressed circumstances.
- In that regard, several market participants have questioned whether the current tight FX bid-offer spreads fully reflect market risk. The spreads may widen rapidly if HFT firms quickly exit the market in periods of market stress. Furthermore, since the FX market is already very concentrated, as suggested by the BIS survey data, any further reduction in the number of key FX participants could lead to a substantial reduction in liquidity during a crisis.
- HFT market-makers have no binding obligation to stay in the market and place quotes in an adverse market environment. HFT-generated liquidity could thus decrease suddenly if market conditions deteriorate considerably, leaving the market with less liquidity just when it needs it most. During the flash crash, there was evidence that some HFT firms did retreat from trading for some time while the volatility peaked, although other participants also retreated.
- To the extent that traditional providers of liquidity have responded to the presence of HFT by being more cautious and selective in their liquidity provision, this may further exacerbate dislocations in market liquidity in times of stress. This would result in a larger movement in prices than was previously observed, even though in normal times the possibility of price gapping is reduced.
- That said, "traditional" market-makers have (and had) no binding obligation to provide liquidity either. Their willingness to do so is similarly driven by profit considerations in the FX business and/or in their client relationships. In the movements in the yen on 17 March, there is evidence that both HFT firms and traditional market-makers pulled back to a similar extent as the yen appreciated sharply. But HFT firms were quick to re-enter the market as the yen reversed.[35]
- HFT has led to these traditional liquidity providers developing proprietary liquidity pools through which to hedge their own risk. They can engage with HFT firms in this

[35] This theme is also evidenced in Chaboud et al (2009).

more controlled environment. However, this takes liquidity out of the main markets such as Reuters and EBS.

6.2 Systemic risks

In general, systemic risk is perhaps more likely to be triggered by a "rogue" algorithmic trade than by pure HFT activity, which tends to involve small trade size, short horizons and diverse strategies.[36] This is certainly consistent with the emerging evidence about the 6 May flash crash in US equities, which seems to have stemmed from a large sell order executed through an algorithm that increased the size of the sell order packages as the market-wide trading volume rose. HFT players initially attempted to provide liquidity against the sell order but the market moved too rapidly, causing them to pull back just as the sell order raised its execution rate in response to the rise in trading volume associated with the initial increase in HFT activity. The resultant reduction in liquidity provision contributed to the price cascade that ensued. This is one illustration of how HFT may not be the trigger of market malfunction but rather an accelerator once the initial problem has been precipitated. That is, HFT has systemic implications because of its potential to amplify and propagate a shock, even if the shock itself does not originate from HFT.

There are, however, some differences in market structure that make a flash crash-type event in foreign exchange less likely than in the stock market during normal trading hours. One fundamental difference is that foreign exchange is a relative price whereas a stock price is an absolute price. Hence while a price movement is a fall for one party, it is a rise for another, thereby redistributing wealth rather than reducing it as occurs when equity prices fall.

Another fundamental difference is that there is a natural demand for foreign exchange related to cross-border trade and financial flows that provides some basic underpinnings to market liquidity. There is no equivalent flow in the equity market.

An issue that a number of market participants raised is that some aspects of foreign exchange trading are becoming increasingly similar to trading in the equity market.[37] Rather than being quote-driven (with market-makers playing a key role in quoting prices and standing ready to meet orders at those prices), the FX market is reportedly becoming more order-driven. Participants are unsure as to whether this is necessarily a bad thing (eg order-driven markets reveal more information about the actual flow of buy and sell orders), but it does potentially lead to the FX market becoming more vulnerable to issues associated with insufficient liquidity. A quote-driven market provides a clearer and more continuous indication of market pricing, which is not present in an order-driven market. That said, an important distinction between the two markets is that the trading environment in equities is regulated whereas foreign exchange is self-regulated by the participants.

6.3 Market integrity and competition

HFT players have often been accused – typically by market participants who have found their margins squeezed – of "abusing" the market by picking off orders as they come to market. As noted above, HFT strategies can actively look for large orders coming to market and quickly take advantage of the effect of such orders on market pricing. To some extent, one could

[36] Examples of systemic risk scenarios in FX: a vacuum of liquidity, a severe dislocation of prices, a failure of a large significant liquidity-providing market participant, and the breakdown of a major FX platform disrupting liquidity and market access.

[37] Market structures are converging under the influence of common driving factors, not just in equity and FX but arguably in all markets.

argue that HFT is just the latest manifestation of the opportunistic trading behaviour that has always been present in the market. The fact that this type of behaviour has been around for a long time, of course, does not make it appropriate market practice.

An issue that might be a cause for concern as a result of this practice is the increase in orders that are processed internally by banks, rather than in the market. The general improvement in technology, rather than HFT per se, has both made the business case for internalisation and facilitated this process. But the success of HFT strategies in moving the market against large orders has further increased the incentive for banks to invest in the technology to internalise more transactions. This internalisation of orders implies that these banks may achieve a certain informational advantage over other market participants that do not see the full flow of orders. Once again, it could be argued that this concern over informational advantage is only a different manifestation of an issue which is ever-present in the market.[38]

HFT is an embodiment (rather than a precipitator) of the operational advantage gained by superior technology. The earlier investors in technology have been able to increase their market share. Once this advantage has been gained, it is difficult for others to compete away this advantage, given the large investment required to replicate the technology. This favours the very large resource-rich firms such as the top-tier FX banks, as well as, in some cases, the small nimble ones such as HFT firms and hedge funds, which can innovate more quickly. It is the middle-tier financial institutions that are disenfranchised.

While, to date, HFT has been most prevalent amongst the major currency pairs, there are clear indications that algorithmic trading, and possibly HFT, are spreading to other currencies, including some emerging market currencies. This again may see a decline in the role of traditional market-makers in those markets as they are unable to keep pace with the technological demands. Price discovery in such markets may tend to shift offshore.

It is important to bear in mind that while HFT is dependent on technology to function and trade positions are often being opened and closed faster than human comprehension, there still remains an important human element. The HFT models are continuously monitored by human traders to ensure that they are performing as expected. In turbulent times, it is a human trader who makes the discretionary call as to whether to turn off the trading model if it has not already ceased trading. Note that the human trader may also switch on trading models that perform well during periods of volatility.

While it may be the case that the impact of HFT on market functioning is generally benign, HFT does require ongoing monitoring to ensure this remains so. The monitoring currently is done by the prime brokers and the trading platforms. There is a concern that PB monitoring may be inadequate if the risk inherent in the services they are providing is not appropriately priced.[39]

In addition to this pricing issue, it is not clear whether the operational ability of the prime brokers to monitor positions is able to keep pace with the speed of HFT. The degree to which this is a problem depends on the speed with which HFT can accumulate an open position before it is detected by the prime broker. The new industry initiatives to help improve PBs' monitoring and management of their clients across trading platforms are a step in the right direction, although it is too early to determine their impact.

Finally, an interesting question to ask, but not addressed here, is whether HFT is a socially efficient use of resources. One can ask this question of a number of fields that adopt cutting edge technology (eg Formula One racing or America's Cup sailing) where the social benefits

---

[38] Banks arguably have the same informational advantage when they receive orders bilaterally via the telephone.

[39] As noted in Section 1.4, this is an issue for prime broking more generally and is not specific to HFT.

are not always immediately apparent. It is not clear that the quest for trading speed, in and of itself, is socially beneficial. But perhaps the more pertinent question to ask is, instead, whether HFT is harmful. One of the purposes of this study has been to provide a set of facts with which one can try to answer that question.

6.4 Looking ahead

One key issue for the future of HFT in FX is how the various ongoing regulatory reform initiatives, especially new regulations on the clearing of OTC derivatives associated with the Dodd-Frank Act (2010), will affect the structure of the FX market. Although spot transactions, which account for the majority of HFT in FX, are exempt from mandatory clearing in the United States, there are still questions at the time of writing as to whether market participants will nevertheless migrate their FX transactions, including spot trades, to venues that capture *all* FX derivatives trading for ease of execution, as well as to minimise margin, capital and collateral costs for non-spot FX transactions. Such a move could have a profound impact on single-bank trading platforms, if these platforms are not eventually designated as a swap execution facility (SEF) under Dodd-Frank.[40] Single-bank platforms have been used increasingly by banks to internalise trading flows and avoid HFT trading. Multi-bank and inter-dealer platforms are expected to be designated as SEFs. The impact of this on HFT participants depends on whether this will lead to more formal regulation of the venues they currently favour and/or whether these participants will face some kind of registration requirement. It also remains to be seen how the real-time public reporting foreseen by the Dodd-Frank Act will affect the behaviour of market participants.

6.5 Concluding remarks

The growth in algorithmic and high-frequency trading in foreign exchange is likely to have significant implications for both the structure and the functioning of the global FX market. Policymakers will need to keep abreast of changes in this space. In some cases, this is happening through the involvement of policymakers in the Foreign Exchange Committees in various jurisdictions. Beyond that, policymakers should develop appropriate contacts and maintain a dialogue with (i) the various foreign exchange trading platforms, (ii) the prime brokerage service providers, and (iii) the algorithmic and high-frequency trading community in FX in order to track developments and to identify key policy issues in a timely fashion.

[40] An SEF is a type of regulated platform on which certain types of OTC instruments ("swaps" as defined by the legislation) will be required by Dodd-Frank to trade. The basic definition in title VII of the legislation is "a trading system or platform in which multiple participants have the ability to execute or trade swaps by accepting bids and offers made by multiple participants in the facility or system, through any means of interstate commerce, including any trading facility, that (a) facilitates the execution of swaps between persons; and (b) is not a designated contract market". Rule writing by regulators (still ongoing at the time of writing) will provide a more exact practical definition, and will thus determine which types of existing trading platforms qualify as SEFs.

Appendix: Empirical literature on algorithmic trading and HFT in equities

Algorithmic trading in equity markets

The larger literature on algorithmic trading in *equity markets* also finds an excess correlation of algorithmic orders and a decline in volatility. However, it suggests that price discovery increasingly takes place through algorithmic activity. Hendershott et al (2011) find that algorithmic trading on the New York Stock Exchange has increased the information content of orders as opposed to trades. Brogaard (2010) and Riordan and Storkenmaier (2010) provide similar results for the NYSE and the Deutsche Börse, respectively. Hendershott et al (2011) argue that since algorithms have virtually no marginal cost in monitoring public information and placing or updating orders, changes in the efficient price will be increasingly revealed in algorithmic quoting activity rather than in trading activity.

HFT in equity markets

Brogaard (2010), using transaction-level data for HFT in US equity markets, finds that:

- HFT participation increases with the market capitalisation of a stock and with order book depth, consistent with anecdotal evidence that HFT firms are primarily active in the most liquid FX pairs and crosses;
- HFT firms employ market-making and contrarian strategies that bet on price reversals; however, there is no evidence of anticipatory trading;
- HFT firms provide the best bid or offer during 65% of the time for large stocks. They consume liquidity in about as many instances as they provide liquidity; and
- HFT firms provide less order book depth than non-HFT participants and are more strategic in liquidity provision.

Regarding total order book depth, findings are inconclusive. Hasbrouck and Saar (2011) find that HFT on the Nasdaq increased order book depth, but Hendershott et al (2011) show that algorithmic trading on the NYSE marginally reduced depth. Hasbrouck and Saar (2011) further highlight that intense high-frequency order submission activity is often abrupt: it will suddenly appear, last for a few seconds, and then quickly disappear. They find that, despite the high number of submissions and cancellations of orders during these episodes, trading activity remains unaffected.

Both Brogaard (2010) and Hasbrouck and Saar (2011) also investigate stressful episodes during the financial crisis and find that the positive impact on market quality generally persists. Brogaard (2010) shows that liquidity provision by HFT firms is negatively related to volatility, but that the effect is economically small. Even on the most volatile days, HFT firms did not stop providing liquidity. However, their consumption of liquidity increased, possibly due to arbitrage opportunities or in order to unload positions in their market-maker function.

References

Brogaard, J (2010): "High frequency trading and its impact on market quality", Northwestern University, mimeo.

Chaboud, A, B Chiquoine, E Hjalmarsson and C Vega (2009): "Rise of the machines: algorithmic trading in the foreign exchange market", Board of Governors of the Federal Reserve System, *International Finance Discussion Papers*, no 980.

Chlistalla, M (2011): "High-frequency trading: better than its reputation?", *Deutsche Bank Research Briefing*, 7 February.

Clark, C (2010): "Controlling risk in a lightning-speed trading environment", *Chicago Fed Letter*, no 272, March.

Commodity Futures Trading Commission–Securities and Exchange Commission (2010): "Findings regarding the market events of May 6, 2010", Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.

Easley, D, M López de Prado and M O'Hara (2011): "The microstructure of the 'flash crash': flow toxicity, liquidity crashes and the probability of informed trading", *Journal of Portfolio Management*, vol 37, no 2, pp 118–28.

Haldane, A (2011): "Race to zero", speech at the International Economic Association Sixteenth World Congress, Beijing, 8 July.

Hasbrouck, J (1991): "Measuring the information content of stock trades", *Journal of Finance*, vol 46, no 1, pp 179–207.

Hasbrouck, J and G Saar (2011): "Low-latency trading", Stern School of Business, mimeo.

Hendershott, T, C Jones and A Menkveld (2011): "Does algorithmic trading improve liquidity?", *Journal of Finance*, vol 66, no 1, pp 1–34.

International Organization of Securities Commissions (2011): "Regulatory issues raised by the impact of technological changes on market integrity and efficiency", consultation report by the Technical Committee of the IOSCO, CR02/11, July.

King, M and D Rime (2010): "The $4 trillion question: what explains FX growth since the 2007 survey?", *BIS Quarterly Review*, December, pp 27–42.

Kirilenko, A, A Kyle, M Samadi and T Tuzun (2010): "The flash crash: the impact of high-frequency trading on an electronic market", mimeo.

Lee, S (2010): "High-frequency trading in FX: open for business", *Aite Group Impact Note*, April.

Riordan, R and A Storkenmaier (2011): "Latency, liquidity and price discovery", Karlsruhe Institute of Technology, mimeo.

Risk, H (2011): "Banks take fight to the algo traders", *Euromoney*, May, pp 72–6.

Securities and Exchange Commission (2010): "Concept release on equity market structure", Release no 34-61358.

Glossary

Aggregator	Trading platform which aggregates prices from multiple sources into one place. It is used by banks, high-frequency proprietary traders and buy-side desks seeking minimal market impact.
Algorithm	A set of instructions for carrying out a procedure or solving a problem in a finite number of steps.
Algorithmic trading	The use of computers and advanced mathematical models to make decisions about the timing, price and quantity of a market order. It is widely used by banks, hedge funds, mutual funds and pension funds. Large trades are broken down into smaller ones to reduce market impact and risk. Trades can be made without human intervention using information received electronically.
Application programming interface (API)	A set of rules and specifications which software programmes follow to communicate with each other; an interface between different software programmes that facilitates their interaction.
Co-location	A practice made possible by exchanges or trading venues renting out space in close physical proximity to their trading servers. Trading firms can use this possibility to place their computers close to the exchange's systems to reduce the travel distance for trading signals and thus latency time.
Dark pools	Anonymous trading venues where quotes on securities or FX are not displayed publicly and trades are executed anonymously.
Electronic communications network (ECN)	Electronic trading system that automatically matches buy and sell orders at specified prices.
Financial Information Exchange (FIX) protocol	An open, public-domain specification developed specifically for the real-time electronic exchange of financial transactions; a messaging standard for the automated trading of financial instruments.
Front-running	The (illegal) practice of using insider information to trade ahead of customer orders. In the context of HFT, however, this term is often used to refer to HFT firms using their mathematical models to detect order flow in order to identify small profit opportunities and using their speed advantage to trade before slower participants can act, although not necessarily using non-public information.
Millisecond	One thousandth (1/1000) of one second.
Latency	The delay between the transmission of information from a source and the reception of the information at a destination. One specific example is the time that elapses between the placement of an order on an electronic trading system and the execution of that order. Can be affected by factors such as geographical distance or bandwidth congestion.
Order book	A list recording the interest of buyers and sellers (bids and offers) in a particular instrument. By observing the order book, one can determine whether the number of bids outweighs the number of offers and use such data to inform one's trading strategy.

Glossary (continued)

Order imbalance	A situation resulting from an excess of buy or sell orders for a specific security in a trading venue, making it impossible to match the buyers' and sellers' orders.
Prime brokerage	A service or practice that enables a bank's customer to conduct foreign exchange transactions in the name of the bank (the prime broker). The prime broker sets up an arrangement that permits the customer to trade directly with dealers in the name of the prime broker. These dealers recognise the prime broker (not the customer) as the counterparty in these trades.
Quote stuffing	The practice of sending multiple orders in a short period of time with no real intention to transact. Orders are typically cancelled by the originator almost immediately. This tactic can be applied in the defensive mode to mask the originator's own trades or to slow down its competitors. It can also be applied in an offensive mode to fake an interest to trade, just to lure others into a deal, thereby exposing their market positioning and scanning the status of the order book.
Single-bank (proprietary trading) platform	A platform developed by a bank internally both for in-house use and sometimes for the use of other banks and non-bank clients on a "white label" or prime brokerage basis. These platforms differ from multi-bank dealing systems in that the primary liquidity provider is a single bank.
Smart order routing	A service that helps traders seek out where prices are best across a range of competing exchanges, platforms and dark pools, and route their orders accordingly for execution. Aggregators use smart order routing.
Top of book	The best bid and ask prices at any given moment.

Members of the Study Group

Reserve Bank of Australia	Guy Debelle (Chair)
	James Whitelaw
Bank of Canada	Harri Vikstedt
European Central Bank	Ivan Fréchard
Bank of France	Sophie Perez
Deutsche Bundesbank	Rafael Zajonz
Hong Kong Monetary Authority	Kitty Lai
Bank of Italy	Maria Lucia Marras
Bank of Japan	Atsushi Takeuchi
Bank of Korea	Hoseok Jung
Bank of Mexico	Alfredo Sordo Janeiro
Monetary Authority of Singapore	Kelvin Ng
Swiss National Bank	Thomas Maag
Bank of England	James O'Connor
Federal Reserve Bank of New York	Anna Nordstrom
Bank for International Settlements	Corrinne Ho (Secretary)

Significant contributions have been made also by Anna Pomeranets and William Barker (Bank of Canada), Jérémy Charbonneau (Bank of France), Gyeong-Ho Kwon (Bank of Korea), Rodrigo Cano (Bank of Mexico) and Jamie Pfeifer (Federal Reserve Bank of New York).

The assistance of contacts at EBS with proprietary data and other information is also gratefully acknowledged.