

Guy Debelle: High frequency trading in foreign exchange markets

Address by Mr Guy Debelle, Assistant Governor (Financial Markets) of the Reserve Bank of Australia, to the ACI High Frequency Trading Conference, Sydney, 12 October 2011.

* * *

Thanks to the participants of the BIS Working Group who all made significant contributions to the report, and James O'Connor and James Whitelaw for helpful comments.

High Frequency Trading (HFT) is somewhat of the *bête noire* of financial markets at the moment. It has been blamed for many adverse market developments in recent times, most famously the “flash crash” in equity markets on 6 May 2010.

But the assessment of HFT is often hampered by difficulties in identifying it, because it is hard to distinguish from other types of automated (but not high frequency) trading, generally known as algorithmic or algo trading. HFT is a part of algo trading but algo trading has a many times greater presence in the market than HFT. 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).

My talk today is based on a report recently published by the Markets Committee of the BIS.¹ The report presents the results of a fact-finding exercise conducted by a study group that I chaired consisting of FX market experts from 14 central banks. One of the primary motivations for the report was that much of the attention thus far has been on HFT in equity markets. There had been relatively little analysis of HFT in foreign exchange, despite its large and growing presence in the market. As a result we felt it was useful to document some of the facts and issues around HFT in foreign exchange. The exercise in itself was invaluable. Beyond the information that we gathered, it helped to strengthen, and in some cases establish, contacts with a number of the main HFT players in the market.

Furthermore, given the different nature, structure and size of the FX market compared to 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.

So how might we think about HFT?

Is the very fact that it is too fast for human perception a major cause for concern?

Let me answer that question with an analogy to another form of high frequency technology.

A BMW is a High Frequency system. There are large numbers of electronic messages flying round a BMW at all times, much faster than you, the driver, can possibly comprehend. You are generally unaware of their presence.

But that doesn't cause you any great concern.

Why? Because you are confident about the engineering and you are confident about the monitoring systems. You are confident that this High Frequency system has been appropriately stress-tested and that the systems are resilient to the stress.

If, on the other hand, you stick a BMW engine in a Leyland P76 and let it loose on the autobahn, there would be a high probability it would crash, potentially taking out some of the surrounding cars with it.

So speed per se is not the problem. The issue is the resilience of the HFT system and the adequacy of the monitoring. So it is useful to think about HFT in foreign exchange in terms of

¹ Bank for International Settlements (2011), “High-Frequency Trading in the Foreign Exchange Market”, BIS Report submitted by a Study Group established by the Markets Committee, September.

these metrics. On resilience, the HFT firm itself has a large incentive to make sure the system being used is resilient. It is typically the firm's own capital at stake. Moreover, there is generally not a lot of leverage involved with HFT. On the adequacy of the monitoring, that is a function of the trading platforms, and particularly the prime brokers. I will return to these issues shortly.

So in my talk today, I will summarise some of the main findings of the Market Committee's report. First, I will describe the topography of the HFT market in foreign exchange and the relationships between the main participants. I will then discuss the effect of HFT on price discovery and liquidity. I will briefly examine the behaviour of HFT in two recent episodes of volatile market conditions, namely the flash crash and the spike in the yen just prior to the intervention in March earlier this year. I will also highlight the key similarities and differences between HFT in FX and HFT in equities.

Finally, I will conclude with some of the lessons learned so far and highlight some of the issues for further consideration.

Market structure

One can usefully think of a topography of the foreign exchange market which starts off at the top with electronic trading and voice trading. The former completely dominates the marketplace these days. Electronic trading can then be divided into manual, where instructions are executed by humans on an electronic trading platform, and automated, where instructions are executed by computer algorithms and there is little or no human intervention. The split here is about fifty-fifty. Mind you, the landscape is changing at a high frequency!

In the FX market, segmenting types of traders is difficult; there is inevitably a lot of crossover in styles and blurring of boundaries. So for the purposes of our study, we split algorithmic (algo) trading into two categories: algo execution and algo decision-making. The former is when a trader uses an algo to execute an order, often used for large orders. The latter is where a model is used to initiate a trade based on parameters such as order book imbalance, momentum, correlations and systematic responses to economic news. It is in the latter camp that HFT is located.

HFT firms generate their revenues from doing a large number of small-size, small-profit trades. They operate with low latency and their risk-holding period is very short, frequently less than one second. As speed is of the essence, indeed the defining characteristic of HFT, co-location is important, that is locating the server that sends the order message as close as possible to the servers of the trading venue. However, our discussion with a number of participants in the market suggested that diminishing returns to speed have well and truly set in, such that the value of future gains may not be worth the investment cost. Indeed, there was a general sense that HFT was reaching a mature phase, where greater returns were more likely to be had from moving into new market segments rather than spending more on enhancing speed.²

HFT participants in FX tend to be located in three cities: Chicago, New York and London, although the physical location of the server is more important than the physical location of the office. The bulk of trading volume appears to be accounted for by several large firms. Most of the activity occurs on the inter-dealer electronic broking platforms (Reuters and EBS) and the multi-bank electronic communication networks (ECNs, such as Currenex, Hotspot and FXall). The various platforms have differing technologies, trading rules and trading parameters. For example, there are varying restrictions on quote life and fill ratios.

² This assumes, of course, that technology is not developed to let one trade faster than the speed of light! Here's where we might get to in the future: "Trading firm hits speed of light", Financial News, 1 April 2011.

To operate in the market, a relationship with a prime broker (PB) is required. There is a handful of large PBs in the foreign exchange market, generally the large investment banks. Our conversations with the market indicated that the terms of access varied a reasonable amount across PBs.

As the PB provides the credit to the HFT firm, it needs the capacity to measure and monitor this credit. HFT firms tend to trade in small size and hold risk only for a very short period, so the outstanding credit at any point in time is likely to be small, particularly compared to other PB clients. But the high-speed nature of HFT can mean that positions can accumulate rapidly, meaning that the PBs need to have rapid monitoring technology to manage their risk effectively. One reflection of this is the recent development of “kill switch” technology to cut off clients’ access to multiple markets simultaneously if unacceptable positions are being accumulated. Andy Haldane highlights the risk of this, given that trading is occurring faster than human comprehension.³

One concern that the report highlights in this regard is whether PBs are adequately pricing the risk in the provision of their services. This risk not only relates to the credit risk but also the potentially detrimental effect the PB customer may have on other parts of the bank’s business, either through spread compression or through reputational risk as the customer typically is trading in the PB’s name. Note that this issue applies to prime broking generally, not just to HFT in foreign exchange.

How large a presence is HFT in the foreign exchange market?

It turns out this is a difficult question to answer. Most surveys of FX turnover data do not separately identify HFT activity. As many HFT firms access the market through their prime brokers, the dealers who report the turnover data will only record the trade as being with the broker, rather than the HFT firm. While Reuters and EBS cannot easily directly identify HFT, an estimate can be obtained by summing volumes of counterparties known to participate in HFT.

Nevertheless, the BIS Triennial Survey of foreign exchange activity provides some indirect evidence on the increased presence of HFT.⁴ Between April 2007 and April 2010, daily average FX turnover grew by almost 20 per cent. Three-quarters of this growth was in spot turnover, primarily in the UK and to a lesser extent, the US. The growth in the UK was principally on electronic broking platforms, while that in the US was mostly on ECNs. Given the electronic broking platforms Reuters and EBS are UK-based and the ECNs, mostly US-based, this is consistent with HFT being a significant contributor to the increase in turnover, given their preference to trade in these marketplaces.

In the report, we calculate an upper-bound estimate for HFT. This estimate is based on the fact that around 45 per cent of trading on EBS and Reuters (depending on the currency pair) is automated. Given that electronic trading itself is around 50 per cent of spot foreign exchange trading, this would imply that HFT could be around 25–30 per cent of all foreign exchange turnover.⁵

³ Haldane A (2011), “The Race to Zero”, Speech at the International Economic Association Sixteenth World Congress, “Approaches to the Evolving World Economy”, Beijing, 8 July.

⁴ See the 2010 BIS Triennial Survey on Foreign Exchange Market Activity and particularly M King and D Rime (2010), “The \$4 Trillion Question: What Explains FX Growth Since the 2007 Survey?”, BIS Quarterly Review, December, pp 27–42.

⁵ The increased market presence of HFT and the impact it is having on the market structure does, however, highlight the need to reassess the way that we are measuring foreign exchange turnover in the BIS survey.

Effect on price discovery and liquidity

What has been the effect of HFT on the market?

The most obvious impact is that spreads are tighter in normal times, particularly at the top of the book. There are questions about the quality of liquidity, however, with claims that liquidity has been impaired for larger orders. Moreover, as we note in the report, the quality of a bid or ask is determined by more than just its price. The size of the quote and its longevity matters too.

The effect of spread compression and depth of book, however, is mostly an issue of redistribution of profit around the marketplace. A number of complaints levelled at HFT firms often reflect the fact that they are taking margin from the party complaining. One example of this is the accusation of predatory pricing. This is sometimes mislabelled front-running. But front-running implies advance “insider” knowledge of a trade. HFT firms don’t have advanced knowledge of the trade as a whole. Rather they detect patterns in trading which are indicative of a large trade coming to market and take advantage of that piece of information (which is available to the whole marketplace). In doing so, they are taking profit at the expense of another participant, but this should be best viewed primarily as redistribution.

There are concerns, however, that while liquidity is improved in normal times, HFT is causing changes in the ecology of the market which result in a worsening of liquidity in stressed circumstances. One of these changes is that some banks are internalising more transactions, which means that less trades are ultimately seeing the public marketplace. A related issue is that traditional market makers may pull back in the normal times as the business becomes less profitable meaning that there may be less liquidity available than in the past in stressed times.

In that regard, we tried to gain some perspective on this issue by looking at the behaviour of the foreign exchange market around the time of the flash crash in equities in May 2010 and the sharp movements in the yen just prior to the coordinated intervention in March this year. The general sense of the evidence from May 2010 is that HFT participants were very active in the foreign exchange market before, during and after the time of the flash crash. However, while they were present, it is an open question about the quality of the liquidity being provided. This episode also supports the proposition that while not necessarily being the initiator of the shock, HFT can propagate a shock brought on by a rogue or poorly specified (non-HFT) algorithm.

In the case of the yen episode, there are indications that HFT players withdrew from the market around that time, but so did others, including traditional market makers. Some of those market makers who remained in the market widened their bid ask spreads dramatically thereby again calling into question the quality of the liquidity being provided. This is standard behaviour, highlighting the point that no participant has an obligation to always provide a quote.

Turning to the issue of the comparison between HFT in equities and HFT in foreign exchange, there are some significant differences between the two. An important difference is the nature of the underlying demand in foreign exchange versus equity. There is wider diversity of participants in foreign exchange and arguably more underlying demand. But there are also signs of convergence, with the foreign exchange market becoming more order driven, rather than quote driven.

Lessons learned and unresolved issues

Having covered some of the facts about HFT in foreign exchange, I will conclude by summarising the lessons and issues that the report highlights:

- On 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 FX market, improving efficiency, and has narrowed spreads, at least for smaller trade sizes. But liquidity for larger trade sizes may have become inferior.

- 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 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 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.
- For example, HFT has led to traditional liquidity providers developing proprietary liquidity pools through which to hedge their own risk. Banks are internalising more transactions in-house using the same technology that makes HFT viable. However, this takes liquidity out of the main markets such as Reuters and EBS. As a result, less information is flowing to the market on trading activity; a development that might detract from the price discovery process, at least for some market participants.
- On HFT and Systemic risks: The 6 May 2010 flash crash in equities shows that rather than HFT per se, algorithmic execution more generally can be a trigger of systemic risk. Rather than being an initiator of such shocks, HFT may accelerate and propagate shocks initiated elsewhere.
- It is important to bear in mind that, while HFT is dependent on technology to function and that trade positions are often being open and closed faster than human comprehension, there still remains an important human element. The HFT models are continually 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.
- On market integrity and competition: Many of the "predatory" or "unfair" practices attributed to HFT, in 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.
- Turning to regulation, 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 we document in the report is whether the prime brokers are technologically capable of keeping 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 varies across platforms. Thirdly, the Foreign Exchange Committees in a number of jurisdictions are currently considering implementing enhanced codes of conduct which aim to address the market integrity issues raised by the increased presence of HFT.

Finally, an interesting question to ask, but not addressed in the report, is whether HFT is a socially efficient use of resources. One can ask this question of a number of fields which adopt cutting-edge technology, such as Formula 1 racing or America's Cup sailing, where the social benefits are not always immediately apparent, but have led to technological

improvements that have benefited society. It is not clear that the quest for trading speed, in and of itself, is socially beneficial. But perhaps the more pertinent way to pose the question is to ask instead, whether HFT is harmful. One of the purposes of the Markets Committee's report has been to provide a set of facts with which one can try to answer that question.

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. It is causing behavioural changes in other market participants that may have an impact on the resilience of the system as a whole.

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 the case. The monitoring currently is done predominantly 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.

Policymakers are endeavouring to keep abreast of the impact of HFT on the FX market by maintaining contact with the different and evolving market participants. In some cases, this is happening through the involvement of policymakers in the Foreign Exchange Committees in various jurisdictions. Industry bodies, such as the ACI, also have a role to play through their work on market codes of conduct.