



Ian Wylie discovers how stock markets are being taken over by super-quick computers and the dangers that could bring

The rise of the machines

UltraDIMM may sound like a damning insult, but the term is the talk of Silicon Valley and describes the latest generation of memory chip that can respond within the space of just five microseconds.

An irrelevance if all we need from our PCs is word processing or refreshing a Facebook page. But it means a great deal – literally – in the world of high-frequency trading (HFT) where microseconds in response time can be the difference between making and losing money on a trade.

High-frequency traders use complex algorithms and fast-responding technology to exploit micro differences in trading prices over time – not in years, months or days, but millionths of a second. And three Warwick Business School (WBS) academics are in the vanguard of exploring not the

facilitating technology but the underlying economic questions of whether the benefits of this fast-growing, automated trading outweigh the costs.

“Almost every market we trade in is turning dramatically automated, to the extent that almost three out of four trades will be run by a machine,” says Vikas Raman, Assistant Professor of Finance at WBS.

Some of us may prefer the outdated image of trading floors in the City or Wall Street where traders shout, wave and clinch deals in a raucous cacophony that concludes with the sounding of a closing bell. But, in reality, the overwhelming majority of stock trades are transacted electronically, and up to 70 per cent of all transactions are now the result of HFT.

So increasingly, the big winners in this new world of

automated stock market trading are not just those with the smartest algorithms but also those with the fastest hardware that can access the information needed to feed those algorithms. News of an event may speed from the wire to the computer network of an ‘algo-trader’ in thousandths of a second. But the difference between recognising and reacting to that data in millionths of a second can mean billions of pounds lost or gained.

It is, as one financial regulator has described it, a ‘race to zero’ as traders go to extraordinary lengths to improve their speeds of reaction, from buying dedicated internet cabling to moving their entire computer networks to be physically closer to the data centres of stock exchanges and news outlets.

Advocates of HFT claim that the automated approach makes for highly efficient and liquid markets, matching buyers and sellers almost instantaneously, while decoupling investors from their dependence on traditional stock exchanges, with the resulting windfall of lower commission prices and other transactions costs.

That view has been supported by three North American academics – Jonathan Brogaard, of the University of Washington, Terrence Hendershott, of the University of California at Berkeley, and Ryan Riordan, of the University of Ontario Institute of Technology, whose paper published by the European Central Bank concluded that high-frequency traders improve price setting and help reduce ‘noise’, or short-term volatility in markets. It warns regulators that introducing measures to curb their activities “could result in less efficient markets”.

Flash crash

However, HFT is cited by some as a major cause of market crashes and creator of volatility. Most notably, many believe the automated withdrawal of liquidity by HFT algorithms accelerated the 2010 so-called ‘flash crash’, where the Dow-Jones Industrial Average plummeted almost 1,000 points before going on to recover most of the losses in just 30 minutes.

And last August, the finger of suspicion pointed again at HFT when the NASDAQ exchange froze for three hours, and Europe’s largest derivatives market shut down for an hour.

According to Raman, HFT has significant implications for small retail investors. “These are not facilities that individuals or even most institutions can invest in,” he explains. “So the advent of HFT makes it very hard for individual retailer investors to be competitive. Secondly, most of the machine trading that happens is done by just a few banks or a few funds. HFT activity is heavily concentrated and it’s not uncommon to go to a modern market and find that 10-15 firms are accounting for 60-70 per cent of trading.”

Likewise, while individual ‘human’ traders cannot trade in multiple markets simultaneously, machines can concentrate on many markets at the same moment and make a market

in different commodities at the same point in time.

“Machines reduce attention costs which means markets are becoming more integrated,” says Raman. “That means inefficiencies and price discrepancies are removed quickly, but because markets are so tightly bound, if there is a shock, contagion will happen at a much faster pace – and that’s what regulators will find very difficult to manage.

“We need to acknowledge that HFT comes with benefits and costs and be cognizant of what are the incremental costs we are paying because of all the wonderful things we are getting from technology.

“HFT makes markets more liquid and more efficient. But there is no such thing as a free lunch. We are getting that liquidity and efficiency at a cost and that cost may be that markets are becoming more fragile”

“There is clear empirical evidence that HFT makes markets more liquid and more efficient. But there is no such thing as a free lunch. We are getting that liquidity and efficiency at a cost and that cost may be that markets are becoming more fragile.”

Because they can react faster than anyone else, Raman’s research suggests algo-traders are the first people to leave a market when something goes wrong. “Our research suggests that markets with high levels of HFT activity can go from lots of liquidity to a dearth of liquidity in a very short space of time. HFT amplifies the impact.”

Algo-traders run

Raman has co-authored a working paper, with Pradeep Yadav, of University of Oklahoma, and Michael Robe, of American University, Washington DC, on fragility in the US crude oil market, which examines the three months prior to September 2006, just before the introduction of electronic trading and when transactions were still done ‘manually’ in pits, against a three-month period in 2011.

“There was a huge influx of automatic traders into the crude oil market after 2006 and our initial analysis suggests that markets behaved differently in periods of distress in the automated era as opposed to how they behaved in the manual era,” says Raman. “Locals increase their participation in times of market stress, whereas machines tend to withdraw and the effect is even more pronounced if the distress is long-lasting. Even during the 2008 financial crisis, the automated trading in the crude oil market reduced dramatically. Algo-traders are the markets’ friends when everything goes well, but they run to the hills when times are bad.”

For Raman’s colleague, Roman Kozhan, the fascination with high-frequency trading began in 2008 when working on a one-month project with a foreign exchange (FX) trader at an investment bank. “On the first morning I sat next to him in his office while he showed me briefly how to trade,” recalls Kozhan.

“And then he said, ‘let’s go and drink coffee, the computer is doing it for me’. So we went and drank coffee, and only returned to his desk for about 15 minutes at 12 o’clock to ‘make sure nothing bad is happening around a macro announcement’. Then we went off again to drink coffee and talk more about our project. I was really impressed!”



▲
The Trading Floor, New York Stock Exchange, 1963

The dark side

Kozhan's latest piece of research looks at the relationship between toxic arbitrage and high-frequency trading in FX. Arbitrage opportunities arise when new information affects the price of one security because dealers in other related securities are slow to update their quotes. These opportunities are toxic since they can result in a trading loss for liquidity suppliers with stale quotes. With Thierry Foucault, of HEC School of Management in Paris, and Wing Wah Tham, of Erasmus University Rotterdam, Kozhan developed a measure of dealers' exposure to toxic arbitrage trades. Using data on high-frequency triangular arbitrage opportunities in the FX market, they show that an increase in dealers' exposure to toxic arbitrage trades can significantly increase trading costs. The finding suggests a possible harmful effect of high-frequency arbitrage activities.

"We are demonstrating that there is a dark side, and that we shouldn't take it for granted that HFT is always good for market liquidity," says Kozhan.

Slicing and dicing

Chen Yao, of Warwick Business School, Maureen O'Hara, of Cornell University, and Mao Ye, of University of Illinois, published research revealing that HFT was leading to more and more big US trades being 'sliced and diced' to less than 100 shares, so that they remained hidden, creating a 'two-tier market'.

They found that a significant portion of US stock trades now occur in odd-lot sizes. For decades, the trading data distributed by exchanges in the US to the public has excluded trades of fewer than 100 shares – known as 'odd lots'. These trades have been considered irrelevant because they were typically made by retail investors, small players who were buying only a few shares at a time.

But as automated trading has grown, more sophisticated

traders have begun using odd lots. "Dividing a round lot into multiple trades may be the result of firms seeking to avoid reporting requirements and may come from those with more information," says Yao.

All of this matters because trades move prices. The fact that savvier traders are using odd lots in large numbers makes that trade activity far more relevant. Following this research, US regulators have now decided that all odd lots will be included in the publicly available data so they will no longer be hidden.

Yao's latest work with Ye focuses on the relationship between HFT and 'tick sizes', the smallest increment (tick) by which the price of stocks can move. Their research argues that a one penny minimum tick size for all stocks priced above \$1 (US Securities and Exchange Commission rule) encourages high-frequency trading. However, they found that non-high-frequency traders (non-HFTers) are 2.7 times more likely than HFTers to provide best prices. Their findings suggest that recent proposals to increase tick size will not improve liquidity. Instead, they will generate further speed competition.

"Liquidity providers earn a profit by providing liquidity. Tick size constrains people from providing liquidity at better prices. For people who provide liquidity at the same price, speed determines who earns that profit," explains Yao.

HFT has become a high-profile target in recent months. In Europe there is mounting pressure from the European parliament to put measures in place to curb HFT. Italy has already begun imposing a tax of 0.02 per cent on many order changes and cancellations that occur within 0.5 seconds of the original order.

The main question now is which, if any, regulatory measures will prove effective. "Regulators do not have a clear picture about what is going on in the market – economic reasons are needed to explain the phenomenon of HFT," says Yao. "HFT is under heated debate, and our paper provides an economic mechanism underlying the phenomenon." ■