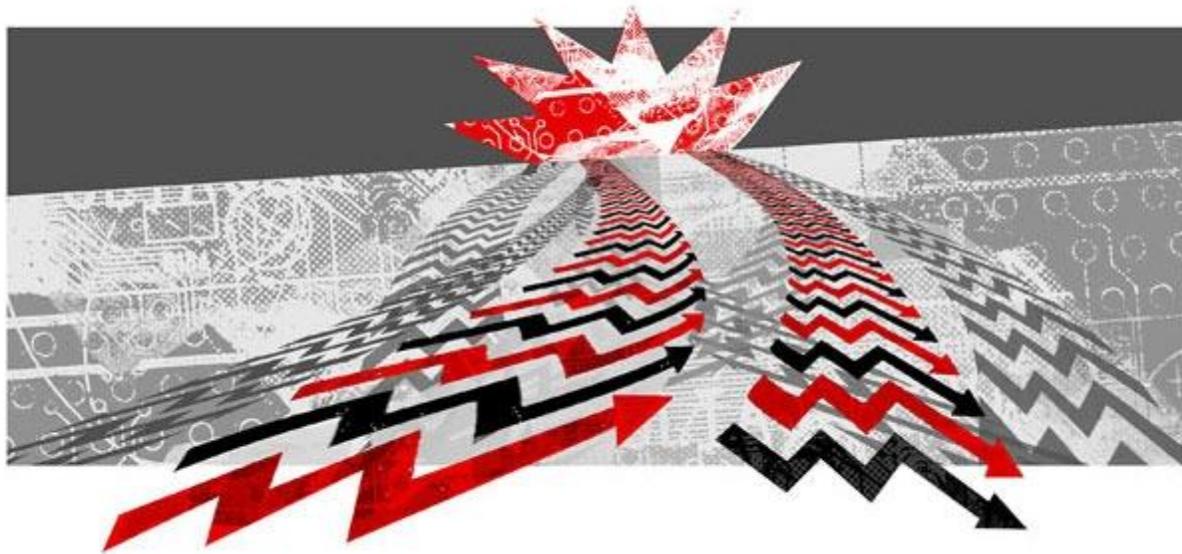


High-frequency trading

The fast and the furious

High-frequency trading seems scary, but what does the evidence show?

Feb 25th 2012 | from the print edition



Gary Neill

ON FEBRUARY 3RD 2010, at 1.26.28 pm, an automated trading system operated by a high-frequency trader (HFT) called Infinium Capital Management malfunctioned. Over the next three seconds it entered 6,767 individual orders to buy light sweet crude oil futures on the New York Mercantile Exchange (NYMEX), which is run by the Chicago Mercantile Exchange (CME). Enough of those orders were filled to send the market jolting upwards.

A NYMEX business-conduct panel investigated what happened that day. In November 2011 it published a list of Infinium's alleged risk-management failures and fined the firm \$350,000. Infinium itself neither admits nor denies any violation of the exchange's rules. It takes the same line on a \$500,000 fine it was given at the same time for alleged transgressions on the CME itself in 2009.

In this special report

- [Playing with fire](#)
- [The ferment of finance](#)
- [The little guy](#)
- [From vanilla to rocky road](#)
- [»The fast and the furious](#)
- [Of plumbing and promises](#)
- [On the side of the angels](#)

- [Safety first](#)
 - [Sources & acknowledgements](#)
 - [Reprints](#)
-

Related topics

- [World markets](#)
- [American markets](#)
- [Industries](#)
- [Investment services](#)
- [Securities Services](#)

Those alleged failures pull back the curtain on some of the safeguards that are meant to protect traders, exchanges and markets from erratic ultra-fast algorithms. The NYMEX panel found that Infinium had finished writing the algorithm only the day before it introduced it to the market, and had tested it for only a couple of hours in a simulated trading environment to see how it would perform. The firm's normal testing processes take six to eight weeks. When the algorithm started its frenetic buying spree, the measures designed to shut it down automatically did not work. One was supposed to turn the system off if a maximum order size was breached, but because the machine was placing lots of small orders rather than a single big one the shut-down was not triggered. The other measure was meant to prevent Infinium from selling or buying more than a certain number of contracts, but because of an error in the way the rogue algorithm had been written, this, too, failed to spot a problem. To complete the catalogue of errors, the firm then allegedly breached another CME rule when an employee used a colleague's trading ID to put on positions that would offset its unwanted exposures.

This incident was unusual for ending in a fine, but in other respects it was not that uncommon. The "flash crash" of May 6th 2010, when American equity markets nosedived by almost 10% in the course of a few nerve-shredding minutes, grabbed popular attention. Although it was not directly triggered by high-frequency traders, the official reports suggested they helped fuel the uncontrolled selling. But there are miniature versions of such flash crashes happening all the time, says John Bates, the chief technology officer of Progress Software, a software firm.

Often they result from algorithms interacting with each other and forming an infinite loop. For a simplified example, take two algorithms that are both programmed always to outbid the best offer in the market, so they will go on outbidding each other. Usually such loops are spotted and stopped, sometimes manually and sometimes automatically, without many people noticing. But the fact that they happen at all feeds the perception that today's equity markets have turned into something more akin to science fiction than a device for the efficient allocation of capital. How, the critics ask, can anyone assess the fundamentals of a company in a fraction of a second? How can lumbering institutional investors, let alone the little guy, hope to compete with the HFTs? And what is to stop a new set of algorithms from unleashing havoc?

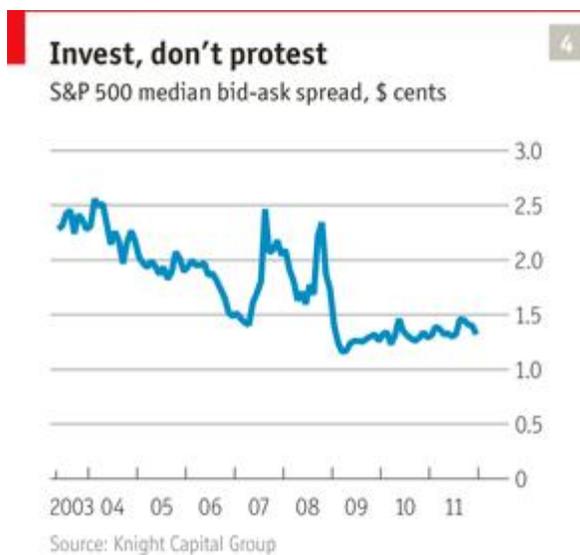
Science v friction

Such questions have gradually drawn the high-frequency traders out into the open. Until recently they saw little need to engage with the wider world. HFTs do not have clients but operate with their own capital. Proprietary algorithms provide a competitive edge, so secrecy is engrained in the culture. But as regulators, politicians and the media focus ever

more closely on their activities, the traders have formed groups on both sides of the Atlantic to represent their interests.

Many are frustrated by what they perceive as an unfair onslaught. "The gap between reality and rhetoric is bigger for this industry than for any other I have seen," says Remco Lenterman, the chairman of the European Principal Traders Association and the managing director of IMC, an electronic marketmaking firm. Plenty of academics support the HFTs' arguments. On the other side are some big institutional investors who accuse HFTs of front-running their orders, and high-profile critics like Bart Chilton, a member of America's Commodities Futures Trading Commission, who has punningly dubbed HFTs "cheetah traders".

To sift through the arguments on both sides is to confront a basic problem with any financial innovation: the difficulty of measuring its benefits. For one thing, there are questions of definition. HFTs are not the only people using algorithms to trade: institutional investors use them to break large orders into small parcels so that markets do not move against them as they execute the order. And although high-frequency trading always involves very brief holding periods and very active trading, it breaks down into lots of different strategies. Some HFTs are momentum traders, riding the wave of a particular trend. Others arbitrage price differences. Others still are marketmakers providing liquidity to buyers and sellers.



Another problem is that there are not enough good data. The fiercest debates centre on the role of HFTs as marketmakers. The evidence tends to favour the HFTs, which can point to a solid body of academic research that shows they increase liquidity, as measured by tighter bid-ask spreads (see chart 4). HFTs also point to testimony delivered to the Securities and Exchange Commission in 2010 by George Sauter of Vanguard, a big fund manager, who concluded that "high-frequency traders provide liquidity and 'knit' together our increasingly fragmented marketplace, resulting in tighter spreads that benefit all investors."

But others say that the increase in trading activity brought about by HFTs, in Europe at least, means that fund managers have to pay additional costs for storage and electronic reporting in order to comply with best-execution requirements. It is hard to disentangle the effects of HFTs on transaction costs from other changes, such as competition among exchanges. A bigger problem, says Paul Squires, the head of trading for AXA Investment

Managers, is that increased liquidity can be illusory. "You can press the button to buy Vodafone, say, and have it executed in a second but in that period 75% of the liquidity has disappeared and the price has moved."

It is certainly true that HFTs are constantly sending and cancelling orders. Some of that activity may be tied to a manipulative technique called "quote-stuffing", in which a flood of orders and cancellations causes congestion on networks and thereby a fleeting trading advantage. But the legitimate explanation for it is that marketmakers cannot afford to be static in case the market moves against them, and that in an ever-faster market HFTs have to be quicker to adjust prices.

"We have got to get away from the idea that speed equals danger," says Richard Gorelick, the "G" in a Texan HFT called RGM. "Professional traders trade continuously and are exposed to market movements all the time, so being able to adjust that exposure quickly gives them confidence to quote better prices." That is why the idea of imposing minimum resting times for quotes before they can be cancelled would almost certainly lead HFTs to widen bid-ask spreads, increasing costs to investors.

To sift through the arguments on both sides is to confront a basic problem with any financial innovation: the difficulty of measuring its benefits

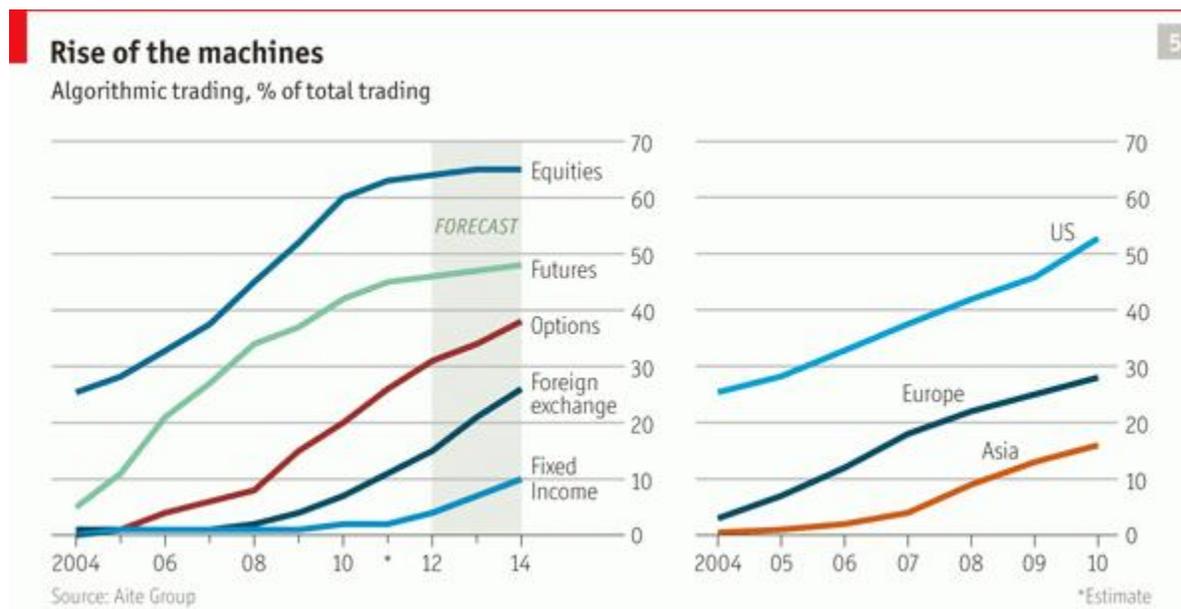
In their 2011 NBER paper Messrs Lerner and Tufano argued that it is virtually impossible to quantify the social impact of a financial innovation because finance involves so many externalities, often unintended ones. For example, it would be almost impossible to measure the aggregate costs and benefits of a fundamental innovation like a bank. Instead, they reckoned, a thought experiment—imagining what the world would look like without a particular innovation—might help.

A world without HFTs is both easy and very difficult to imagine. Easy, because the old world of specialist marketmakers and floor trading existed only a few years ago, so people remember it well. There is little obvious enthusiasm for returning to that model. Not only were transaction costs higher but the same arguments about unfair advantages were being put forward in different forms. Now the complaints are about the milliseconds HFTs gain over ordinary investors by putting their servers right next to the exchanges' data centres; then they were about the monopolistic privileges of the specialists and the advantages of being on the floor. Institutional investors may complain about being forced into "dark pools" (off-exchange venues where they can deal anonymously) to avoid HFTs, but these pools existed before HFTs and were set up in part to avoid being scalped by brokers or floor traders.

But imagining a world without high-frequency trading is also hard. That is because the HFTs are what Larry Tabb of TABB Group, a research firm, describes as an "outcrop" of the market structure. They are a natural outcome in a world in which trading is automated, and in which there is competition between lots of different exchanges and a need for someone speedily to knit together the prices they offer. "The real question is whether humans make worse mistakes when they write algorithms or when they trade," says Terrence Hendershott of the Haas School of Business at the University of California, Berkeley.

In practical terms, trading history is highly unlikely to be reversed. Regulators in developed countries have no evidence that radical change is needed, nor any appetite for it. Developing countries, which have become the standard-bearers of sensible financial

regulation, are racing towards high-frequency trading as they seek to improve liquidity (see chart 5).



Meanwhile the industry itself pushes inexorably forward. That certainly entails greater speed: the industry used to think in terms of milliseconds (it takes you 300-400 of these to blink) but is now fast moving to microseconds, or millionths of a second. It also means smarter algorithms. People have gone from trading in open-outcry pits to trading via screens to programming algorithms. The next stage, says Mr Bates of Progress Software, will be self-learning systems, in which sentient algorithms mine the capital markets, spotting correlations that are too complex for humans to see and suggesting trading ideas as a result. Humans will still be needed to validate these ideas, he says reassuringly.

But the prospect of even faster markets raises the problem posed by the Infinium case and by flash crashes large and small: the threat from HFTs to the stability of markets. Software has a nasty habit of developing bugs. Algorithms behave unpredictably once they are out of a testing environment and into the market proper. Even leaving aside the potential for deliberate market abuse, traders will sometimes get things wrong.

Plugging the holes

So real-time risk-management safeguards have to be put in place that work at different levels—of the HFTs, the prime brokers, the exchanges and the regulatory agencies. Some of these safeguards already exist. The exchanges have limits on orders above a certain size, for instance, and on the number of orders that can come in. But there are always holes to be filled.

The May 2010 flash crash revealed a glaring problem with the structure of American equity markets, for instance: that circuit-breakers which automatically pause trading if there are violent price swings kicked in only once the entire market reached a certain threshold. The regulators have since introduced “single-stock circuit-breakers” for any stock that moves up or down by 10% or more in a five-minute period. Most observers are confident that this would stop another flash crash in its tracks.

However, a focus on equity markets may distract attention from other asset classes where HFTs are present and growing. The Bank for International Settlements, a club of central bankers, last September issued a useful fact-finding report on the role of high-frequency trading in the vast foreign-exchange markets. It worried about the capacity of the prime brokers, through which HFTs gain access to credit, to keep pace with their clients. If they do not, the prime brokers themselves are exposed to the possibility of HFTs rapidly accumulating risky positions.

Again, the industry is aware of this risk. Traiana, a post-trade processing firm, launched a software program last year that aggregates clients' positions across lots of different venues in real time and activates a "kill switch" that stops clients trading once pre-defined limits are breached. But this is a voluntary initiative, not one that is required by regulators. Other markets remain unprotected. Regulators should not be afraid to act firmly to define and enforce standards for market surveillance and trading controls across venues and asset classes. And getting the infrastructure right is important in other areas too.

<http://www.economist.com/node/21547988>