Principal® Global Equities

RESEARCH • RESOURCES • RESULTS

Rise of the Machines

On the Evolution of High-Frequency Trading

Contributing Author: Huw Gronow – Senior Equities Trader – Principal Global Equities

Trading in equities evolved into a formalized environment in 1611 with the opening of the Amsterdam Stock Exchange. The building, designed by Hendrick de Keyser, was the first structure created specifically to house the trading of shares – in this case, shares of the Dutch East India Company. Even though the physical structure of the Amsterdam Stock Exchange fell into disrepair and was demolished in 1835, it wasn't until the early 1970s that the true human-to-human nature of financial markets began to dissolve. Most significantly, on February 8, 1971, the NASDAQ stock exchange began trading – the world's first electronic stock market. While no more than a glorified electronic bulletin board at its inception, the NASDAQ did help to begin lowering trading costs and pave the way for what was to come. Nine months later, Intel would publicly introduce the first microprocessor. The Intel 4004 was revolutionary in that it combined all of the elements of a computer's central processing unit onto a single chip. It also delivered, in the size of a fingernail, the processing power of computers that had taken up entire rooms. And in the way that microprocessors would supplant large transistors soldered to circuit boards, so too would electronic trading begin to supplant the traditional open-outcry system. Fast forward to the twenty-first century, trading and technology have advanced to such a degree that trades can be executed in fractions of a second at the cost of thousandths of a cent per share. This has allowed for the rise of a type of trading that makes the most of technology, that moves faster than humans can perceive, that puts machines in control of billions of dollars, and that has divided the public perception of institutional finance – high-frequency trading.

This paper will examine the conditions, both regulatory and technological, that laid the groundwork for high-frequency trading. It will describe the good and the bad of high-frequency trading in an attempt to give the reader an idea of the potential benefits this oft-maligned practice can provide, as well as its hazards.

High-frequency trading, or HFT in industry parlance, did not spring forth fully formed from the mind of some genius trader. HFT required three initial conditions for its growth: a plethora of trading venues, a dramatic increase in the speed of trading, and microscopic trading costs. A combination of regulatory changes and technological advancement was the impetus for the accelerated evolution of HFT.

New Regulation Starts the Ball Rolling

In the old regime, incumbent trading venues, such as the New York Stock Exchange (NYSE)[1] in the United States, dominated the market in a virtual monopoly of trading activity. Since they existed essentially alone, the monopoly exchanges could charge whatever they liked and had little incentive to reduce the costs of trading for investors. Two pieces of regulation began to change all of that. In 2005, directives in the United States (Regulation National Market System) and

[1]The Amsterdam Stock Exchange eventually combined with other European exchanges to become Euronext, which was itself subsumed by the NYSE Group in 2007 to create NYSE Euronext.

in Europe (Markets in Financial Instruments Directive) set out to foster competition and benefit the end investor. The result of the new regulations, implemented a year or two later, was the upset of the dominant, incumbent exchanges.

What Competition has Wrought

A Plethora of Trading Venues

The new regulations worked as they were intended. Almost total domination by incumbent exchanges was supplanted by a new, relatively equal distribution of business among many smaller exchanges. The tide of new exchanges brought forth a structural shift in the trading landscape. Through the liberalization of concentration rules, where the incumbent exchanges enjoyed, for the most part, a monopoly, new types of trading venues sprang forth in direct, aggressive competition. Traders could develop strategies that would route to what became known as "lit" or "dark" exchanges. Lit exchanges are those where the available liquidity is publicly displayed; investors see the prices at which participants are willing to trade. Dark exchanges give an investor the ability to cloak themselves in anonymity to help avoid unintentionally exposing their trading strategies; however, the cost of this privacy is a lack of assurance of execution.

U.S. Equity Market Share Breakdown - H1 2012

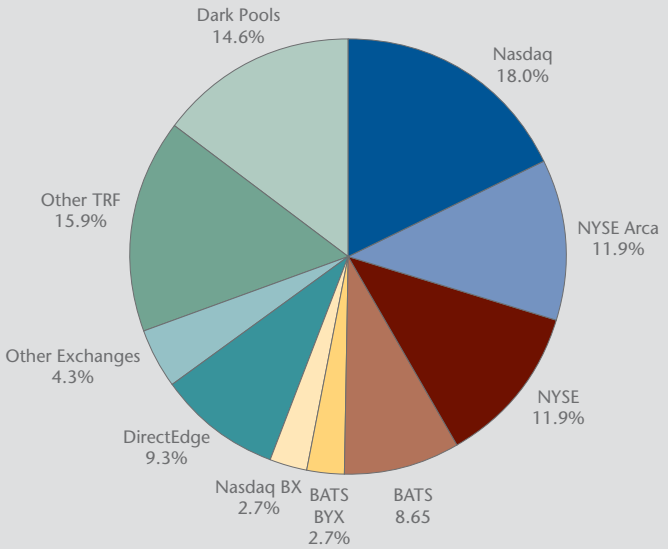

"Other TRF" includes broker capital commitments and internalizations
"Other Exchanges" includes NYSE Amex, CBSX, Chicago, Nation and PSX Date for 2012-H1
Source: Credit Suisse Trading Strategy

Speed

This new competition also brought about an increase in the speed of trading. The new trading venues that sprouted in the United States all had a common feature: fast, electronic access to trading. In the earliest days of open-outcry trading, the speed of the markets was limited by the speed at which two human beings could interact. Speed has always been important, but then computers took over part of the equation, and the speed of the markets was then limited by how fast humans could interact with the machine that was executing the trade. Now, with computers integrated at every point in the trading process, the speed of the market is limited only by how fast one computer can interact with another computer. Given its new electronic nature, trading is now essentially limited only by the speed of light – technically, the speed of light through a fiber-optic cable.

As an example, one trading-communications firm operates a fiber-optic network between New York and Chicago to provide a highly reliable, ultra-fast, electronic connection between Wall Street and the HFT firms that sprang up in the Midwest. The distance between New York and Chicago is 825 miles. The speed of light through a vacuum is approximately 186,000 miles per second. Through fiber-optic cable, light moves at around 125,000 miles per second. This means that the round-trip latency (the time elapsed between trade placement and execution) of a signal between Chicago and New York is approximately 13.33 milliseconds. To put that in context, assume your local grocery store had an algorithm that allowed you to do all of your daily shopping in 13.33 milliseconds. At those speeds, you could do a year's worth of shopping in four seconds.

That said, the next stage in the race for speed is already beginning – taking out the physical limitations of data bouncing around in cables and being able to approach the speed of light in a vacuum – via microwave wireless data networks.

As a single point of comparison, one such wireless network delivers a round trip in 155 microseconds, around 35% quicker than through a fiber optic cable over the same distance. That is, 0.000155 seconds – a very quick shopping trip indeed.

The abundance of new trading venues results in a need to split large trades into smaller orders to try to maximize the probability of success in execution, thus further shrouding an investor's trading strategy and potentially limiting the price movements that large trade could engender. Increasing the speed at which trading takes place makes this a highly competitive race to capture liquidity. However, humans cannot process all the variables required to profitably pull off the operation; algorithms, or computerized math-based trading instructions, are needed to facilitate what has become known as smart order routing. Say an investor wishes to purchase one million shares of stock; it would be foolish to push that order in its entirety onto a lit exchange. Like tipping your hand at poker, you would essentially tell the other market participants what you're trying to do and they would move against you. The solution is for the algorithm to break the large "parent" order into smaller "child" orders. The so-called "smart order router" serves to direct those trades through the various venues to see if the liquidity can be acquired at the appropriate premium and in the appropriate size.

Microscopic Trading Costs

Quick and frequent trading can become extremely expensive if the cost of trading set out by exchanges doesn't decrease. By using a pricing model based on volume, the new trading venues were able to lead customers away from the incumbent exchanges. The maker-taker model arose as the primary way of providing liquidity[2] in the new exchanges. In the maker-taker model, margins shrink past "razor thin" to become "microscopic" – 30 millionths of a cent per 100 shares traded (0.3 cents per share). The providers of liquidity, or market makers, are incentivized to trade because they are paid that tiny fraction of a cent to do so. The takers of liquidity (usually institutions) are charged a slightly higher amount to effect the trade, and the exchange keeps the difference. Since trading only a few times a day would generate only a few millionths of a cent in profit, the exchange has to see tens or hundreds of millions of trades every day to make the model profitable. Now we find ourselves in the realm of a new type of market maker. One who is willing to provide or remove liquidity on the bleeding edge of this new technological infrastructure that combines incredibly high speeds with infinitesimally low costs and an abundance of exchanges. Meet the high-frequency trader.

High-Frequency Trading

It's a mistake to define high-frequency trading as a single type of strategy. Broad, sweeping generalizations of HFT lead to blanket statements like "high-frequency trading is bad" or "high-frequency trading hurts investors." Firms that operate HFT techniques are neither good nor evil, HFT is merely a set of techniques, or a description of a set of tools; and like all tools, it has the potential for misuse. Moreover, it is important to remember that these firms trade for themselves; they do not represent underlying clients or investors. But given their reliance on these techniques, these firms are referred to as "HFT firms". HFT is a subset of algorithmic trading

High Frequency Trading Estimates

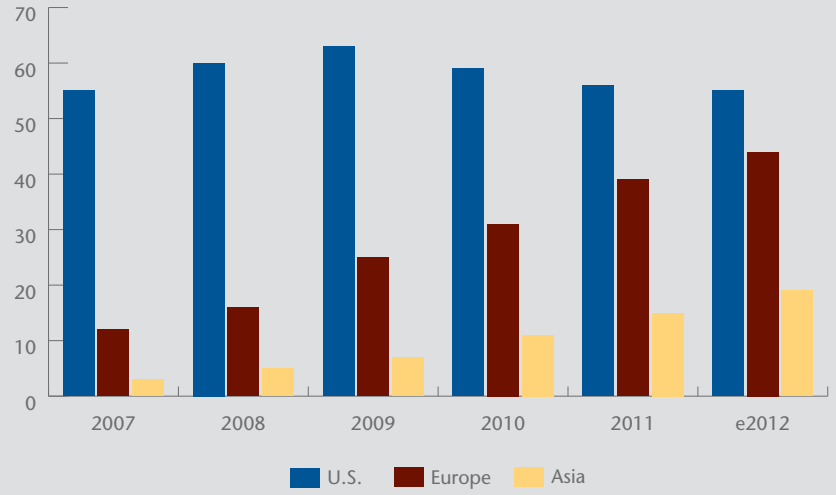

Source: Aite Group

[2]"Providing liquidity": this can be perceived (due to the use of the word "providing") as passive only. Contemporary reality is that this in fact refers to any order/quote arriving to the market, be it passive, or indeed, aggressive.

and has been loosely described as an automated trading process that would lose its efficacy if even one second of delay were added to the trading process. Said another way, the shorter the trading horizon is, the more effective HFT becomes in terms of short-term alpha – with extraordinary Sharpe ratios (that is to say almost risk-free returns).

Types of HFT strategies

Dividing HFT strategies into two camps based upon the outcomes they pursue gives us both the beneficial and the insidious. Legitimate, helpful high-frequency trading strategies include both market-making activity and arbitrage strategies. Importantly, however, they differ greatly from the market makers of old, who would provide liquidity in good times and bad, arguably providing inherent stability – one of the important aspects to consider following the Flash Crash of 2010. In market-making activities, high-frequency traders provide liquidity to the markets, and have the effect of tightening spreads at the top of the book – but again, not providing "depth" of liquidity. This is said to improve market efficiency. Arbitrage strategies are those that look for minute aberrations in the prices of different instruments on different exchanges, then seek to make money by closing out those discrepancies or by eliminating them – all at blistering speed. These strategies work to tighten prices in the market.

In a certain sense, the more sinister aspects of HFT are no new story. As long as there has been trading, there have been those who seek to illicitly benefit at the expense of others. There are laws against this type of market behavior and we strongly feel that these existing laws should be enforced and activities like insider trading, "trading ahead of", and general market manipulation should be eradicated. With HFT strategies that are viewed as negative, markets face the same villain as always; however, now moving at incomprehensible speed, he cannot be easily detected. In less than the blink of an eye, like some lightning-fast Fagin, disreputable HFT strategies are able to analyze, execute, and profit – meaning that regulators face a major challenge in surveilling the market for new, or even old, forms of market abuse.

There are three general types of HFT strategies that can be considered to have no benefit to the end investor: market manipulation, momentum ignition, and latency arbitrage. Market manipulation strategies run the gamut from front running (trading ahead of a known client order to profit from the bump that order will produce) to quote stuffing (flooding the market with quotes to slow it down and examine the movements) and layering (putting in spurious bids and orders to give the false impression of the market's position at any one time). While these strategies don't rely exclusively on HFT, they are made all the more dangerous because of the camouflage that such speed provides.

With momentum-ignition strategies, the high-frequency trader takes a position, and then attempts to move the market in a direction to support that position. One way this can work is by using the HFT algorithms to rapidly submit and cancel orders in an attempt to "fool" the algorithms of other traders and push prices away from fundamental values, and ultimately toward the intended position.

One might also point to so-called "rebate arbitrage" (i.e., "non-profitable") strategies as having an unintended effect. Pricing incentives for liquidity providers was, as we have already seen, intended to foster competition and remove the incumbent exchanges' hegemony. However, when a strategy looks to earn both the rebate for posting and the rebate for taking, the effect can be to artificially inflate volumes and attract volume triggers to participation-based strategies on arguably a false premise, similar to the momentum ignition strategies.

One last category of shady HFT practices is the latency arbitrage strategy. Essentially, this is the process of exploiting the small delays in price dissemination by taking advantage of structural inefficiencies. The unfair advantage comes from HFT firms buying access to the raw data feed from the market, which gives them an informational advantage over most market participants, who receive the consolidated data feed. Then, the HFT firm will reduce the physical distance between their own computers and those of the exchange. This allows the HFT firms to back into the consolidated data feed before the rest of the market receives it, then trade ahead of the rest of the market based on their analysis. While neither access to raw data feeds, nor co-locating servers are illegal, the overall effect is one of a manipulative and disreputable nature. This strategy is brought even further into question by allegations that the NYSE was purposely providing data to some HFT firms before other market participants. While the NYSE was

Speed of Light – Distance = Trading Advantage

In 2010, NYSE Euronext opened a new data center in Mahwah, New Jersey. Anticipating the needs of high-frequency traders, NYSE Euronext sold over 40,000 square feet of co-location space to HFT firms. In a business where speed is everything, this co-location effectively puts the HFT firms mere feet away from the NYSE's new matching engines, providing them with the lowest available latency[3].

[3]Co-location is an extremely precise process. Even minute differences in cable length can result in microsecond differences in latency.

fined, it neither admitted nor denied the charges. Though, the mere hint that an exchange could be complicit in this type of activity harms investor confidence.

In fairness, the SEC is working on it. Recently, they found that LeveL, an alternative trading system (ATS) in the US run by a firm called eBX, told clients that their flow of orders to buy or sell securities would be kept confidential and not shared outside of LeveL. However, LeveL was in fact allowing an outside technology firm to receive and store information about LeveL subscribers' unexecuted orders for its own business purposes. That firm's routing business therefore received an information advantage over other LeveL subscribers. The SEC charged eBX with failing to protect confidential trading information of users and issued them an $800,000 fine.

Our View

The conditions that make high-frequency trading a possibility – speed, low costs, and a multitude of trading platforms – have been an unarguable success for investors. Saving both time and money has been of almost incalculable benefit to investors as trading costs have plummeted to fractions of fractions of a cent and speed of execution has approached the speed of light. The question is whether the potential negative effects of high-frequency trading outweigh the benefits that have amassed.

An event like the "Flash Crash" of May 6, 2010 was an especially dramatic example of the volatility that high-frequency trading can sew. In the United States, the Dow Jones Industrial Average tumbled 990 points in approximately five minutes. It then snapped back just as abruptly, recouping 543 points in less than two minutes. Much of this volatility was due to HFT algorithms reacting, then overreacting, to incoming market data. Because of the number of HFT strategies in the market, this can become like a negative feedback loop – selling that forces algorithms to sell even more.

The most recent and alarming event was undoubtedly the Knight Securities fiasco which appears essentially to be an example of surveillance shortfall, both regulatory and internal. The latest debate is whether a "kill switch" can effectively be deployed to avoid another Knight-like event which for them, without the injection of capital from other firms to cover the $440m loss they incurred in just 45 minutes of trading, would likely have seen the immediate demise of one of the US' biggest electronic market making firms.

In our view, the technology itself is not the enemy; computerized trading systems are no more evil than light bulbs or airplanes. Besides, the genie is out of the bottle; there's no conceivable way that markets would allow trading to devolve to some earlier state where execution is measured in minutes rather than milliseconds. Rather, we feel it is the unchecked structural problems of the markets that are eroding the public's confidence. Witness the Facebook IPO debacle, where NASDAQ's systems were too jammed up to process orders as they were received on the first day of trading. Or consider the botched BATS IPO, where BATS shares crashed and briefly traded around a penny on their own exchange, an eerie recall of the Flash Crash of 2010. Now, investors rightly demand compensation, which keeps concerns over market structure firmly on the front pages.

As an analogy, governments and the general public are fine with cars that are able go 140 miles per hour, as long as they're confident that the authorities are able to identify, stop, and prosecute speeders. So it is with the markets, the technology and infrastructure have advanced so quickly that the oversight has failed to keep up. Slowing everyone down isn't the answer. New, overly burdensome regulation is also not the answer. Market manipulation is already a crime. For example, there already exists the EU Market Abuse Directive, which needs to be updated to more accurately reflect newly identified practices like quote stuffing, and new metrics like message-to-trade ratios which at inflated levels must have some empirical distortion on the trades that do end up taking place. Furthermore, what the market needs is a faster, smarter way of identifying these abusive HFT strategies as they happen, as they do for slower miscreants, and a market infrastructure that can better seal off the negative effects to avoid contagion, which in turn restores and reinforces confidence in today's equity markets.

About Principal Global Equities

Principal Global Equities, a specialized investment management group within Principal Global Investors provides client-focused investment solutions spanning equity markets worldwide. We are fundamental investors focused on bottom-up stock selection within a sophisticated comparative framework. Our research universe encompasses over 10,000 companies, large and small, in emerging and developed markets. We provide expertise in active-core and specialty strategies, with risk profiles aligned to distinct client objectives and preferences.

Disclosures

The information in this document has been derived from sources believed to be accurate as of October 2012. Information derived from sources other than Principal Global Investors or its affiliates is believed to be reliable; however we do not independently verify or guarantee its accuracy or validity.

The information in this document contains general information only on investment matters and should not be considered as a comprehensive statement on any matter and should not be relied upon as such. The general information it contains does not take account of any investor's investment objectives, particular needs or financial situation. Nor should it be relied upon in any way as forecast or guarantee of future events regarding a particular investment or the markets in general. All expressions of opinion and predictions in this document are subject to change without notice.

Subject to any contrary provisions of applicable law, no company in the Principal Financial Group nor any of their employees or directors gives any warranty of reliability or accuracy nor accepts any responsibility arising in any other way (including by reason of negligence) for errors or omissions in this document.

All figures shown in this document are in U.S. dollars unless otherwise noted.

This document is issued in:

- The United States by Principal Global Investors, LLC, an advisor registered with the United States Securities and Exchange Commission.
- The United Kingdom by Principal Global Investors (Europe) Limited, Level 4, 10 Gresham Street, London EC2V 7JD, registered in England, No. 03819986, which has approved its contents, and which is authorised and regulated by the Financial Services Authority.
- Singapore by Principal Global Investors (Singapore) Limited (ACRA Reg. No. 199603735H), which is regulated by the Monetary Authority of Singapore. In Singapore this document is directed exclusively at institutional investors [as defined by the Securities and Futures Act (Chapter 289)].
- Hong Kong by Principal Global Investors (Hong Kong) Limited, which is regulated by the Securities and Futures Commission.
- Australia by Principal Global Investors (Australia) Limited (ABN 45 102 488 068, AFS Licence No. 225385), which is regulated by the Australian Securities and Investments Commission.

In the United Kingdom this document is directed exclusively at persons who are eligible counterparties or professional investors (as defined by the rules of the Financial Services Authority). In connection with its management of client portfolios, Principal Global Investors (Europe) Limited may delegate management authority to affiliates that are not authorised and regulated by the Financial Services Authority. In any such case, the client may not benefit from all protections afforded by rules and regulations enacted under the Financial Services and Markets Act 2000.

Principal Global Investors is not a Brazilian financial institution and is not licensed to and does not operate as a financial institution in Brazil. Nothing in this document is, and shall not be considered as, an offer of financial products or services in Brazil.