



GUEST EDITORIAL

Larry Harris, CFA
Fred V. Keenan Professor of Finance
Marshall School of Business, University of Southern California
Los Angeles

What to Do about High-Frequency Trading

Like powerful tools or drugs, high-frequency trading (HFT) is both extraordinarily valuable and incredibly dangerous. Although HFT greatly reduces average trading costs for investors, it also poses systemic risks to the markets, hurts investors through front running, and decreases investor confidence.

Concerns about the potential damage from HFT have produced many high-frequency complainers, especially among buy-side traders. Their concerns are real, and regulators must address them.

But regulators also must be careful. Poorly conceived policies could easily hurt the markets. For example, current proposals to require minimum resting times for orders would damage markets without producing much benefit.

The debate about HFT has been quite emotional, in large part because people naturally fear what they do not understand well.

A seemingly obvious but incorrect argument has also influenced the debate: Many people believe that restrictions on HFT cannot harm the markets because investment decisions are not made over one-second intervals, much less over millisecond intervals. The premise of this argument is right, but the conclusion is wrong. HFT promotes markets by making them more liquid and thus ultimately lowers corporate costs of capital. High-frequency traders need to submit and cancel their orders quickly to provide liquidity cheaply.

The most pressing danger that the markets face from HFT is least recognized: High-frequency traders are engaged in a costly technology arms race that will not end well for investors if regulators do not act soon. Fortunately, a simple change in order-handling procedures—described herein—can sensibly stop this race.

Identifying what regulators should and should not do about HFT requires some understanding of what high-frequency traders do.

A Quick Brief

HFT is trading for which success depends critically on low-latency communications and decision making. High-frequency traders use computers to process electronic data feeds, make trading decisions, and convey orders to electronic exchanges over intervals measured in micro- and milliseconds. HFT has grown with the electronic exchanges that enable it.

HFT strategies vary considerably. Most high-frequency traders use dealing and arbitrage strategies to offer liquidity or to move liquidity among markets. Some high-frequency traders trade on news feeds about fundamental values. Lastly, a few high-frequency traders actively front-run other traders.

Valuable HFT. High-frequency traders who use dealing and arbitrage strategies make markets liquid by providing investors with opportunities to trade. Numerous reliable studies have shown that transaction costs for both retail and institutional traders have decreased substantially with the growth of HFT.

The cost savings are easy to understand. In comparison with human dealers, computers

- have perfect attention spans,
- follow instructions exactly,
- do not allow emotion to cloud their judgment,
- can watch and learn from thousands of sources of information simultaneously,
- do not cheat, and
- cost less and require smaller offices.

These advantages have greatly reduced transaction costs because many high-frequency traders compete with each other to provide liquidity. High-frequency traders have largely displaced traditional human dealers when they compete in the same markets.

Regulators must be very careful that they do not inadvertently harm the high-frequency traders who make markets liquid.

*Guest Editorial is an occasional feature of the Financial Analysts Journal.
This piece reflects the views of the author and does not represent the official views of the FAJ or CFA Institute.*

Harmful HFT. Some high-frequency traders use computers to monitor and interpret electronic news feeds. When they identify material information, they immediately trade on the favored side. These traders cause prices to reflect information about fundamental values faster than the prices otherwise would.

Traders who post standing limit orders that do not yet reflect the changes in value implied by news lose to high-frequency traders. These liquidity-supplying traders include dealers (most of whom are also high-frequency traders) and patient investors. Investors have always lost to better-informed traders, but many resent that they lose simply because they cannot learn about important events as quickly as high-frequency traders can.

The incremental economic benefit of prices made more efficient by seconds is hardly meaningful: Corporations do not make operating decisions that depend on informative stock prices—for example, decisions about capital structure or compensation—in the seconds immediately following news releases.

Regulators should protect investors from these losses by requiring companies to notify exchanges when they expect material information will be revealed during trading hours so that the exchanges can halt trading before the news arrives. Many already do. For the same reason, many governments release major information only when the markets are closed or at pre-announced times, before which liquidity suppliers generally cancel their standing orders.

Very Harmful HFT. A few high-frequency traders front-run buy-side traders who are working orders, thereby making the latter's trades more expensive. Such activities are legal if the high-frequency traders do not improperly obtain information about the orders they front-run.

Regulators would like to stop this type of HFT, but most policies that they might implement would have serious unintended consequences.

The strategies that high-frequency traders use to front-run other traders vary by whether they front-run orders that they expect traders to submit (order anticipation) or standing orders that traders have already posted (quote matching).

Order anticipators examine trades and quotes to detect when traders are using algorithms to split up large orders that will move the market. They then trade ahead of such orders to profit from expected price changes. The successful implementation of this strategy depends less on low-latency communications than on high-quality pattern-recognition algorithms. The order anticipation problem is thus not really an HFT problem.

This strategy is ancient. Like poker players, traders have always watched each other carefully to identify what they might do next. Computers are now essential to the successful implementation of order anticipation strategies because they can often recognize certain patterns faster and more accurately than people can and because much trading is now electronic.

Regulators can do little to protect buy-side traders from order anticipators without delaying or reducing the dissemination of quotes or trade reports. Either solution would make it harder to predict future orders, but both would make markets substantially less transparent, which would hurt investors.

The best hope for protecting large, algorithmic buy-side traders is to reduce the information about trade sizes that markets disseminate. In particular, instead of reporting actual trade sizes, markets should report only approximate trade sizes within various buckets or report only aggregated volumes at intervals of 5–10 minutes. These changes, together with the use of hidden orders and dark pools by large traders, would substantially reduce the ability of traders to identify future orders.

Quote matchers profit by extracting option values from standing limit orders submitted by slower traders. They trade ahead of such orders by improving prices slightly or by trading in another venue. If prices then change in their favor, they profit. But if they expect prices to move against them—perhaps because the prices of correlated securities or indices have changed—they immediately exit their positions by trading with the standing limit orders. The traders who issue the standing orders thus fail to trade when they wish they had and trade when they wish they had not.

Like order anticipation, quote matching has always been a problem for large buy-side traders. It was a primary source of profit for exchange specialists before electronic trading became common.

The success of the quote-matching strategy depends on how quickly traders can (1) cancel their unexecuted orders when the standing orders that they are front-running are canceled or filled and (2) trade with these standing orders when they want to exit their positions before the orders are canceled or filled by other traders. Low-latency high-frequency traders have thus come to dominate this strategy.

Regulators could protect buy-side traders from quote-matching high-frequency traders by prohibiting high-frequency traders from canceling their orders too quickly. A minimum resting time would cause high-frequency traders to lose more often when markets move against them while they are trying to establish their positions.

Unfortunately, this rule—which regulators are actively considering—would also cause liquidity-supplying high-frequency traders to lose more often when offering liquidity, which would ultimately increase investor transaction costs. The harm done to market liquidity would be much greater than the benefit obtained from discouraging quote-matching high-frequency traders because they trade much less often than liquidity-supplying high-frequency traders.

The Technology Arms Race

High-frequency traders are engaged in an arms race.¹ To beat their competitors, each trader is spending increasingly large sums on expensive technologies to speed their trading. If actions are not taken to stop this arms race, investors will be worse off and economic welfare will decline.

High-frequency traders compete by offering better prices or more size when quoting to trade. They also quickly cancel their orders to avoid losing when market conditions change, and they occasionally initiate trades with other traders when attractive opportunities appear.

Being very fast is not enough to be a profitable high-frequency trader. Such traders must be faster than their competitors. The fastest high-frequency traders get the best places in line when quoting to trade, avoid trading when they no longer want to trade, and take valuable trading opportunities when they first arise. Slower high-frequency traders lose because they have to stand in line behind their faster competitors, because those competitors often trade with the slower traders to their disadvantage when market conditions change, and because those competitors take valuable trading opportunities before the slower traders can.

High-frequency traders go to great lengths to be faster than their competitors. They locate their servers next to exchange servers to minimize communication times. They pay for special high-speed data feeds and for the shortest communication lines between exchanges. They use extremely fast computers. They write hyperefficient computer code. And increasingly, they even hard-code their software onto silicon chips to minimize response times.

The long-run implications of this arms race are not yet well appreciated: The fastest high-frequency traders will eventually drive out their slower competitors, and only a few HFT firms offering liquidity—perhaps just one or two—will survive. The high costs of acquiring the technologies needed to be fast enough to compete successfully will become an insurmountable barrier to new competitors. Indeed, these costs already block all

but the most wealthy and wildly optimistic potential competitors.

Decreased competition among high-frequency traders will be a troubling outcome of this winner-take-all arms race. When the competition among high-frequency traders thins out, the remaining traders will no longer have to quote aggressive prices to obtain order flow. Investors will have to pay higher prices when they buy, and they will receive lower prices when they sell. The costs of trading will rise.

Economic welfare will suffer because high trading costs make investing less attractive to investors. Corporations that need to raise capital for new projects will have to sell securities at lower prices to attract investors, which will increase their capital costs. Fewer projects will be funded, and fewer jobs will be created. Everyone will be worse off if this arms race is not stopped.

Fortunately, a small and easily implemented change in exchange rules could substantially reduce the incentives to acquire the expensive trading technologies now required to compete successfully as a high-frequency trader. Regulators should simply require all exchanges to delay the processing of every order instruction they receive by a random period of between 0 and 10 milliseconds.

Without this rule, any high-frequency trader with merely a one-millisecond advantage over a competitor will always beat that competitor. With this rule, the faster high-frequency trader will beat the slower one only 59.5% of the time. (If the two high-frequency traders were equally fast, the rate would be 50%.) Both traders would still want to be faster, but the benefits of speed would be greatly reduced.

This small change would substantially reduce technology expenditures by high-frequency traders without any negative effect on the quality of the markets. Instead, by lowering the costs of entry, it would ensure that HFT remains a highly competitive business in which traders primarily compete against each other by improving prices and quoted sizes. The current competition—in which high-frequency traders invest in technologies whose only benefit is to give them an advantage over their competitors—provides no benefit to public investors.

Systemic Problems

Poorly designed or poorly used electronic trading systems pose systemic risks. In particular, trading systems that demand too much liquidity too quickly can cause prices to fall or rise to unreasonable levels. Many electronic trading systems can also generate so much order flow that they clog order-routing/processing systems and thereby

deny market access to other traders. These problems may arise when

1. an algorithm goes out of control and submits unanticipated orders (Knight Capital),
2. a trader misuses an algorithm by setting parameters that cause it to trade too aggressively (the May 2010 Flash Crash in U.S. equities),
3. the algorithms that traders simultaneously use get into a negative feedback loop in which they take turns responding to each other, or
4. terrorists, anarchists, national enemies, disgruntled employees, or mentally unstable traders obtain and exercise malicious control over an order-generating/routing system.

Although market destabilizing, the first two of these scenarios are self-correcting because they invariably lead to unacceptable losses to the traders ultimately responsible for the unintended orders. For example, the trader who used an algorithm to sell \$4.1 billion of S&P 500 Index e-mini futures contracts (thereby precipitating the Flash Crash) lost about \$150 million of his clients' assets in only 10 minutes. Fears of such losses cause well-run firms to devote substantial resources to avoiding them.

The greater problem for markets lies in algorithmic feedback because it is more likely to lead to excess order traffic that disrupts markets than to large trading losses that traders seek to avoid.

Regardless of their origins, these events all cause external damage to other traders (and to market confidence in general), for which the responsible traders are not penalized. Thus, expected trading losses may not provide sufficient incentive to reduce the incidence of such events, especially for poorly run firms in which those responsible for avoiding trading losses do not fully appreciate all trading risks. Indeed, the number of recent problems caused by algorithms suggests that firms have not paid sufficient attention to these issues.

Regulators should intervene by requiring that all firms that generate electronic orders have a kill switch. This proposal is reasonable: Firms that create orders by identifying various patterns in market data certainly can, and should, examine their own outgoing order flow to identify patterns inconsistent with their business models.

Preventing the malicious use of trading systems requires a different solution because anyone who can control a trading system may also be able

to disable its kill switch. To avoid these problems, all exchanges must monitor their incoming order flow to kill any inappropriate orders, or they must regularly inspect the kill switches to ensure that they have not been tampered with.

Conclusion

The vast majority of high-frequency traders benefit investors by creating more-liquid markets. Despite this well-documented fact, many commentators have called for slowing HFT by imposing minimum standing times for orders. However, faced with this constraint, high-frequency traders would quote less aggressive prices for smaller sizes to avoid losses to better-informed traders. This policy would have the unintended effect of increasing transaction costs for public investors.

Markets need to be slowed, but not because HFT is dangerous. Markets need to be slowed slightly to wisely stop an arms race that will eventually decrease competition among high-frequency traders and thereby increase investor transaction costs. Minimum standing times would not address this problem. Instead, we should delay the processing of all orders by a trivially short, random period of between 0 and 10 milliseconds, which would ensure that high-frequency traders always provide markets with very low transaction costs.

Why do buy-side traders complain so much about HFT? Perhaps because they are often caught between their portfolio managers and the markets in which they trade. All portfolio managers want better executions for their orders. When pressed about a disappointing execution, traders often find it easier to blame the markets than to accept responsibility.

Regulators give much credence to the opinions of buy-side traders on market structure issues because they are expert traders working in the trenches on behalf of public investors. But regulators should also remember that traders have complained about market structures for years. They have often been right, but not always.

HFT has made markets more liquid than ever by substantially reducing the costs of dealing. Regulators must act to protect this liquidity by ensuring that many high-frequency traders will always compete to fill the orders of public investors.

Notes

1. This section borrows from my recent op-ed in the *Financial Times*, "Stop the High-Frequency Trader Arms Race" (27 December 2012).