Does US high-frequency trading need stricter regulatory oversight?

YES

The crux of this question is whether there is anything new or particular to high-frequency trading (HFT) that demands special regulatory attention. Definitions of HFT vary, and the label is imprecise as it groups together a variety of distinct practices and tactics that may affect markets in different ways. Nevertheless, all agree that HFT is characterised by extremely high speed, and autonomous operation. Both of these properties can cause markets to perform in modes not possible before the advent of HFT, and we have substantial evidence of anomalies caused by HFT activity. So what is it that makes HFT qualitatively different, such that it needs regulatory scrutiny?

Speed

First, the very speed at which HFT operates introduces qualitatively new patterns of activity. High-frequency traders receive and process information from exchanges with ultra-low latency measured in milliseconds, and are able to react with orders within a few milliseconds more. Such speed provides an edge in responding to economic news, which translates to profitable advantage as the first responder reaps the benefit of executing against mispriced outstanding orders.

Most of the time, however, there is little substantial news – in the majority of milliseconds, nothing of real-world economic consequence actually happens. As such, the advantage is actually deployed in response to events within the market itself.

Momentary mispricing arises even in the absence of news, simply due to the short-term imbalances of supply and demand induced by order flow. The winner-takes-all nature of response to mispricing leads inevitably to a latency arms race.

This sees competing high-frequency traders take extraordinary measures involving enormous investments in technology and communications infrastructure to shave microseconds off their response time and thereby improve their odds of winning the profit opportunities that arise throughout the trading day.

Autonomy

Second, the autonomy of HFT increases the difficulty of anticipating its effects in complex or unusual situations. Autonomous operation is of course necessary to compete in the latency arms race, and the ability to process high volumes of information at great speed is a great spur to automation.

Human traders may likewise be unpredictable, but we have hundreds of years of experience with them, and the relative slowness of human-scale trading can limit damage accrued before measures are taken.

New behaviours

That HFT may introduce different behaviours to the market is not in itself a bad thing. Indeed, the general progress in automation of trading is quite likely responsible for reductions in transaction costs and improvement in price discovery. What the point about qualitative differences does entail *a priori* is that HFT deserves special scrutiny in case such extremes of speed and autonomy contribute to systemic risks, or allow particular trading tactics that detract from efficiency or otherwise degrade the performance of financial markets.

The *a priori* argument is only strengthened by evidence from experience. The most notorious market HFT anomaly to date was the so-called Flash Crash of May 2010. Whereas some consider HFT to have been exonerated as the proximal cause of that incident, many knowledgeable observers argue that the predominance of HFT actors in the affected markets enabled the subsequent response, by suddenly withdrawing the liquidity they had generally been providing. And although the Flash Crash left no permanent damage and has not been repeated, as far as we understand it could happen again at any time. In fact, it has been reported that so-called mini flash crashes in individual securities have actually been occurring with some regularity in recent months.

> We have substantial evidence of anomalies caused by HFT activity

Flash crashes

The question of flash crashes underscores how little we actually understand about the implications of pervasive HFT activity. Numerous academic studies on HFT in the last few years – based on theoretical models, or empirical data analysis – have shed light on aspects of HFT. But the overall picture remains incomplete at best.

Standard theoretical finance models, generally speaking, abridge the fine granularity of information transmission time scales driving the latency arms race. And even the most thorough data analyses have difficulty identifying the effects of particular HFT strategies. What we care most about – the potential exacerbation of rare and extreme market events – is inherently difficult to study. This is because, by definition, the conditions present themselves only rarely.

All this suggests a need for more fundamental research for the long-term, and continued vigilance in the short-term. Prudence in the face of uncertainty calls for regulatory scrutiny, but also dictates a posture of humility and cautiousness in promulgating specific measures.

The most effective lever

The proper goal of regulatory oversight is not to diminish automated trading *per se*, but rather to promote the beneficial HFT tactics and discourage the practices deemed malign.

Prohibiting the perceived harmful tactics directly is often difficult to enforce, given that trading strategies are generally unobservable. In many cases, the most effective lever is the design of market rules that align incentives, by rendering the discouraged strategies more costly or less effective.

By Michael P Wellman, professor of computer science and engineering at the University of Michigan