Man vs. Machine: Liquidity Provision and Market Fragility

Vikas Raman, Michel A. Robe and Pradeep K. Yadav*

"High-frequency trader liquidity, evident in sharply lower peacetime bid-ask spreads, may be illusory. In wartime, it disappears." – Andrew Haldane, Bank of England, July 2011[1]

1. The Policy Question

The major equity markets of the early nineties – NYSE, NASDAQ, and London Stock Exchange – were dealer markets with "market-makers" affirmatively obliged to provide "liquidity", i.e., always stand ready to buy (or sell) when a public retail or institutional trader wanted to sell (or buy), and thereby maintain robust and orderly markets. In contrast, equity markets are now largely electronic limit-order book markets, and these markets are robustly liquid only when some traders continually fulfill the role of providing abundant liquidity by posting standing buy and sell limit orders for other public traders to execute against. However, these liquidity-providing traders act *entirely voluntarily,* without any affirmative obligations to maintain liquid and orderly markets. They supply liquidity only to earn, rather than pay, the bid-offer spread; and only when it is optimal for them to do so as part of their overall trading activities. Hence, an important area of concern for exchanges, regulators, and market participants is the consistent and continual availability of liquidity – limit orders to execute against – not just in good times but also in periods of turbulence and stress, so that the traders who need liquidity can always reliably get immediate execution of their orders. This concern has been significantly heightened following the rapid growth in algorithmic (i.e., machine) traders (hereafter "AT"), who can move in and out of markets very quickly, and who now overwhelming dominate liquidity provision through computer-based automated trade decision and execution systems that harvest bid-offer spreads without any human trade-by-trade intervention or pre-meditated directional bets, participating on both sides of the book, turning over inventory intraday with minimal capital investment as often as optimal. Regulators and policymakers have extensively voiced concerns that, while AT improve liquidity on average, they also generate greater dangers of periodic episodic illiquidity.[2] In this context, the question of enormous policy relevance is the following: do machine traders run for the exits when the going gets tough, and thereby make markets more fragile in periods of stress? Automated trading by robots is the inevitable march of technology, and an in-depth understanding of the associated systemic risk in liquidity provision is critically germane to determining the optimal regulatory response that can keep today's markets acceptably robust.

---

* Vikas Raman (vikas.raman@wbs.ac.uk) is at the Warwick Business School, University of Warwick, England. Michel Robe (mrobe@american.edu) is at the Kogod School of Business, American University. Pradeep Yadav (pyadav@ou.edu, corresponding author) is at the Price College of Business, University of Oklahoma, 307 W. Brooks Room 205A, Norman, OK 73019, USA. This white paper is adapted from a working paper with the same title and the same authors presented at the NSE/NYU Annual Research Conference in August 2014, and available at https://www.nseindia.com/research/content/1314_BS8.pdf.

[1] Speech at the International Economic Association Sixteenth World Congress by Andrew G Haldane, Executive Director, Financial Stability, Bank of England, available at http://www.bis.org/review/r110720a.pdf.

[2] See, for example, *Foresight: The Future of Computer Trading in Financial Markets (2012) Final Project Report*, The UK Government Office for Science, London, page 11: Executive Summary. Similarly, Mary Schapiro, Chair, U.S. Securities and Exchange Commission, said in September 2010: "…*high frequency trading firms have a tremendous capacity to affect the stability and integrity of the equity markets …. in tough times…*".

The issue of systemic risk in liquidity provision, and consequent market fragility, was brought into sharp focus by the Flash Crash of May 2010, and continues to remain in focus with the regular occurrence of mini-crashes.[3] This paper accordingly presents the results of an empirical investigation of the participation and transactional liquidity provided by AT – the "machine" – relative to human traders – the "man" (hereafter "MT") – during periods of market turbulence or stress, in contrast to what they do in "normal" periods; and the resultant implications for the quality, fragility, and regulation of equity markets.

2. What do we expect?

Periods of market stress and turbulence are characterized by high levels of information intensity, or information asymmetry, or information uncertainty. On one hand, it can be argued that AT have a competitive speed advantage over MT during such periods, since AT can access and process data faster and more efficiently through their pre-programmed artificial intelligence algorithms, without being constrained by limits to human cognition and bounds of human rationality (Biais and Woolley, 2011; Biais, Foucault, and Moinas, 2015).[4] On the other hand, it is argued that algorithms pre-programmed *ex ante* cannot deal with the complexity of turbulent periods as effectively as manual traders. This is because "data" or hard news releases are not necessarily the same as "information" – they have to be processed into usable information. Zigrand, Shin and Beunza (2011) state this eloquently: "*... algorithms know the price of everything and the value of nothing.*"[5] AT should have a clear advantage over MT when the processing is simple: e.g., exploiting any opportunities by comparing prices of related assets, or across fragmented markets. However, pre-programmed algorithms are unlikely to be effective when potentially unrelated and unusual data or news clips have to be processed into economically relevant price information. This is because: (a) artificial intelligence cannot be pre-programmed to comprehensively cover the entire feasible spectrum of preference uncertainties and economic complexities;[6] and (b) periods of market stress are rare and unique, reducing the benefits of rules based on the past, and increasing the risks of serious glitches while running new or adapting old algorithms. Consequently, it is argued that, in turbulent, unpredictable conditions, AT minimize their losses by simply applying the "kill switch" (Yadav, 2015)[7]; and the evidence in Zigrand, Shin and Beunza (2011) supports the widespread use of kill switches in selected periods. An important corollary of this complexity argument is that AT should lose their informational advantage over MT during extreme events/periods of market stress.

---

[3] See, for example, Sornette, Didier, and Susanne Von der Becke (2011), "Crashes and High Frequency Trading", Swiss Finance Institute Research Paper No. 11-63.

[4] See: (a) Biais, Bruno, and Paul Woolley (2011), "High frequency trading." *Manuscript, Toulouse University, IDEI*; and (b) Biais, Bruno, Thierry Foucault, and Sophie Moinas (2015), "Equilibrium Fast Trading" *Journal of Financial Economics* 116 (2): 292–313.

[5] See: Zigrand, Jean-Pierre, Hyun Song Shin, and Daniel Beunza (2011), "Feedback Effects and Changes in the Diversity of Trading Strategies", in *The Future of Computer Trading in Financial Markets - Foresight Driver Review – DR 2*, London.

[6] Preference uncertainties are modelled by: Biais, Bruno, Johan Hombert, and Pierre-Olivier Weill (2014), "Equilibrium Pricing and Trading Volume under Preference Uncertainty", *The Review of Economic Studies.*

[7] See Yadav, Yesha (2015), "How Algorithmic Trading Undermines Efficiency in Capital Markets", *Vanderbilt Law Review*, Volume 68, 101-163.

A second reason why AT may reduce their participation and liquidity supply in turbulent markets relative to MT arises because AT liquidity supply activity is characterized by very limited commitment of capital and ultra-short intraday horizons. AT bear position risks only when they expect to profitably offload their positions within their ultra-short trading horizon. The AT trading advantage stems from their ability to trade in and out of positions faster than others.[8] Such agility is hindered when capital is locked-up in a single position. Therefore, the lower the chances of profitable inventory rebalancing in a short period of time, which will be the case in a one-sided "extreme" market, the greater the reluctance to take a position – and, conditional on participation, the smaller the position undertaken. Furthermore, AT's over-arching imperative of keeping their capital commitment low means that they are much more likely to frequently trade out of positions in turbulent markets by demanding liquidity rather than continuing to function as liquidity suppliers.

We know that the trading of AT is significantly more correlated than that of MT.[9] In this context, the complexity and short horizon arguments articulated above suggest that the behavior in normal and extreme periods should be different, because ATs would consider kill switches in extreme turbulent periods, and the simultaneous application of kill switches by AT traders could arguably lead to a severe deterioration in liquidity. In a stress situation, many algorithms can quickly coordinate and act simultaneously and feed each other, potentially giving rise to feedback loops that make markets fragile. Zigrand, Cliff, and Hendershott (2011) argue that such feedback loops are the underlying force behind financial crises, and are more likely to arise, or at least harder to supervise, in AT environments.[10]

3. What do we do?

In spite of the extensive regulatory concerns, and the two important reasons – greater complexity in turbulent periods and AT's short horizons – that could precipitate swift withdrawal of participation and liquidity supply from AT in turbulent periods, extant empirical research has focused only on "normal" market conditions, and concluded that, on average across all periods, algorithmic/high-frequency trading enhances liquidity and market quality.[11] In contrast, the contribution of this research is to focus on periods of market turbulence and stress, where stress is measured by high and persistent volatility, and/or high and persistent order imbalances, and/or high and persistent bid-ask spreads. We empirically test whether AT participation in trades and the contribution of AT to transactional liquidity supply – i.e., posting of standing buy and sell limit orders that have provided trade execution

[8] See, for example, Jovanovic, Boyan, and Albert J. Menkveld (2015), "Middlemen in Limit-Order Markets." Working Paper, VU University Amsterdam and NYU.
[9] See, for example, Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega (2014), "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market", *Journal of Finance*, 69 (5): 2045–2084.
[10] Zigrand, J., D. Cliff and Terrence Hendershott (2011), "Financial Stability and Computer Based Trading" in *The future of computer trading in financial markets*, Foresight, Government Office for Science, pages 6–23.
[11] See the following: (a) Hasbrouck, Joel, and Gideon Saar (2013), "Low-Latency Trading", *Journal of Financial Markets* 16 (4): 646–679; (b) Hendershott, Terrence, Charles M. Jones and Albert J. Menkveld (2011), "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance* 66 (1): 1–33; and (c) Hendershott, Terrence and Ryan Riordan (2013), "Algorithmic Trading and the Market for Liquidity" *Journal of Financial and Quantitative Analysis,* 48(4): 1001–1024.

and immediacy to other traders – is as reliable and stable as that of MT even in times of market stress. We also test whether this makes equity market liquidity more fragile.

The empirical analyses in this paper are based on trades and orders data from the National Stock Exchange of India ("NSE"). NSE data are particularly suitable for this study because the NSE has always been an electronic limit order book market with voluntary liquidity supply, and algorithmic execution was permitted only after a clearly specified date in 2008 – we exploit this feature by comparing trading before and after algorithmic trading. In contrast, the U.S. equity markets were at least partially dealer markets with affirmatively obliged market makers, and not always purely electronic without any trading floor, which makes it difficult to isolate the specific contribution of the algorithmic nature of voluntary liquidity supply to any observed potential for fragility. Further, our data also provides broad trader classifications and flags algorithmic trades within each classification; thereby, enabling benchmarking algorithmic traders with other manual, voluntary traders of the same trader type. Such a benchmarking provides a cleaner estimate of the effect of automation on trading strategies. Therefore, the NSE provides an excellent laboratory to investigate the impact of AT on market fragility. Our results are based on comparing two periods across different categories of traders: May 2006, during which there were no AT (i.e., all were MT), and May 2012, during which there were both AT and MT.

4. What do we find?

We document several results of considerable regulatory and policy importance, and of direct relevance to public retail and institutional traders.

1. First, we examine liquidity provision in duly consummated trades, as measured by the proportion of AT participation in trades, and the proportion of trades in which AT were actually supplying liquidity in as much as they were the counterparty with the standing limit order in the trade.
    a. We find strong evidence that, in contrast to manual traders, AT significantly reduce their participation and transactional liquidity provision in periods of extreme turbulence as signaled by significantly high and persistent volatility, customer order imbalances, and/or bid ask spreads.
    b. AT participation and liquidity provision in consummated trades drops by about 25% on average for all types of extreme periods.
    c. The withdrawal from liquidity provision in consummated trades is extremely significant and large in magnitude – more than 40% – for AT who are external customers of the exchange; and it is these external customers who are the main suppliers of liquidity.
    d. The overall withdrawal from liquidity provision in consummated trades is small and largely insignificant for AT who are exchange members, but the extent of liquidity supplied by them is much less.

2. Second, we examine limit order activity during periods of stress, measured as the proportion of the number and the volume of new orders submitted by AT and MT traders; and the proportion of the number and volume of net-new orders (new orders minus cancelled orders) submitted by them.
    a. All categories of algorithmic traders, including exchange members, withdraw from the order book significantly more than their manual peers during stressful market conditions.
    b. For algorithmic traders who are external customers of the exchange, the proportion of new orders and net-new orders drops dramatically by 6 and 11 percentage points respectively, which are 8.4 and 6.1 times the corresponding variable for corresponding manual traders. Similarly, the volume of new and net-new orders decreases massively by 4 and 18.1 percentage points, which are 4.7 and 2.7 times the corresponding variable for corresponding manual traders.
    c. Exchange member algorithmic traders withdraw the least, but they also withdraw very significantly more than their corresponding manual peers.

3. Third, we examine the change in the cost of liquidity provision is measured as the change in the relative pricing or aggressiveness of the new algorithmic and manual orders that are actually submitted during periods of market stress.
    a. We find that those algorithmic external traders that remain in the market place and continue posting limit orders in stressful conditions, place significantly less aggressive orders than their manual peers, effectively increasing the price at which they are willing to supply liquidity.
    b. However, there is no significant change in the relative aggressiveness of orders placed by exchange members.

4. Fourth, we find that the propensity of algorithmic traders to withdraw is strongly dependent and conditional on the extent of persistence of abnormal market conditions. This is consistent with the withdrawal of algorithmic traders being related to both the short horizon of such traders, and the potential constraints in the ability of pre-programmed algorithms to deal with the complexity of market signals in turbulent periods.

5. Fifth, we find that the withdrawal of algorithmic traders in stressful markets corresponds one-to-one with a loss of their informational advantage with respect to manual traders in such markets. This indicates that speed based information advantages of algorithmic traders disappear in the complexity of market signals in turbulent periods, and this motivates algorithmic traders to withdraw from voluntary liquidity provision in periods of stress.

6. Finally, we find the withdrawal of algorithmic traders has a significant propensity to generate feedback loops that can make markets more "fragile". Specifically, we find the following:
    a. A reduction in algorithmic trading or algorithmic liquidity provision significantly increases the probability of extreme market conditions. A one standard deviation decrease in algorithmic trading or algorithmic liquidity provision increases the odds of an extreme event by about 30%.
    b. The potential for fragility is further exacerbated by the fact that algorithmic traders in a stock withdraw significantly from that stock even in the absence of stressful conditions in that stock, when another similar sized stock experiences an extreme event. The withdrawal of algorithmic traders hence displays significant contagion and correlation across stocks, even when stressful market conditions do not.

5. Concluding Remarks

Overall, our empirical results indicate that, in contrast to human traders adapting manually in (higher latency) real time, ex-ante pre-programmed algorithmic trade execution is less conducive to low impact adjustment of complex information asymmetries or flows, and increases the systemic risk of robust liquidity provision in turbulent times. Hence, our results concretely reinforce the extensive regulatory hunches that have hitherto existed about the potential for systemic fragility created by algorithmic trading.

That said, algorithmic trading represents the inevitable march of technology that is here to stay. We know that algorithmic trading is beneficial most of the time since it reduces the average cost of trading for both retail and professional traders, and also has a positive impact on market quality. However, our research shows that, at the same time, algorithmic trading also increases the dangers of periodic episodic illiquidity, and it is important for regulators and policy makers to develop optimal policy responses that can significantly mitigate these dangers, while continuing to harvest the benefits of algorithmic trading.

We recommend a policy response in the form of a liquidity enhancement scheme that focuses on creating a voluntary cadre of designated liquidity providers who are selectively and differentially provided financial incentives with two objectives. First, we can incentivize aggressive standing limit orders specifically in periods of stress, defined transparently on a look-back basis in terms of volatility levels, with the incentive being an order of magnitude greater in extreme periods relative to low or medium volatility periods. Second, we can incentivize posting of standing limit orders simultaneously on both sides of the book, with a high threshold for the proportion of time a particular trader is on both sides of the book, so that she has an incentive to continue to remain a liquidity provider in troubled times. We further recommend that this risk of periodic episodic illiquidity should be clearly recognized by the exchange, and the economic cost of this focused liquidity enhancement (to traders signing up as designated liquidity providers) clearly borne by all other traders through an explicit separate fee charged from each trader as a fraction of the value of shares traded.