EasyChair Preprint

№ 3950

Financial Markets in a Digital World: Electronic Trading Systems and Market Fairness

Terry Kennedy

July 26, 2020

Financial Markets in a Digital World: Electronic Trading Systems and Market Fairness

Terry J Kennedy

Abstract

Before the days of automated trading, stock markets were driven primarily by a mixture of fundamental analysis, technical analysis, speculation, and human emotion. Today, much the same is true. However, the prevalence of electronic automation has seen a marked change in way that financial markets behave in today's environment.

This paper looks at finding answers to the debates surrounding the fairness and regulation of Electronic Trading Systems (most notably, High Frequency Trading.) We will begin by building a framework upon which we can base our subsequent analysis before moving on to address the various arguments.

1 INTRODUCTION

As noted by Angel and McCabe (2013), the concept of automated trading itself cannot be labelled as unethical per se, as it is only a tool. So, to be in a position to discuss the merits or pitfalls of these technologies, we must first consider how these tools operate within the context of the financial markets in which they exist. Before we can define the relative worth of any sub-system, or set of sub-systems (such as algorithmic trading) within a given market, it is necessary to ask the following question: "How do we define an effective (or 'healthy') financial market?"

In their paper 'Ethics for Automated Financial Markets' (Cooper et al., 2020) it is argued that in a quest for the definition of an 'effective market', we should be concerned solely with whether or not that market operates in such a way that the socially beneficial outcomes of that market can be realised. Additionally, it is posited that any 'effective market' should necessarily exhibit the following characteristics*:

1. Voluntariness^
2. Transparency
3. Informational Efficiency
4. Reliability

By inference, we could then argue that in order for automated trading systems to have a positive impact on the host market, they should not operate in opposition to these characteristics. If they do so, then they would be contributing negatively to the overall health of the market, and therefore ought to be removed and/or banned. Through this line of reasoning, we can now seek to determine the relative worth of arguments for and against the use of High Frequency Trading.

2 THE RISE OF AUTOMATION

In this paper, we will be focussing on a sub-type of automation called High Frequency Trading (HFT). HFT is a form of algorithmic trading that allows the users to buy or sell huge volumes of a given asset class at speeds which are orders of magnitude faster than could be achieved by any human. In the United States, the market share of High Frequency Trading algorithms has been estimated to be between 40 and 70 percent (Fabozzi et al., 2011).

In more recent years, the accumulation of vast quantities of market data along with leaps in technological efficiency have facilitated the creation of incredibly advanced and highly complex algorithms fundamentally targeted at profiting from financial markets. The ways in which this is achieved varies depending on the particular strategy that is being employed. An important thing to understand about algorithmic trading as a whole, is that the models that underpin these technologies (and the data upon which they rely) are not necessarily limited to one instrument or even one exchange, but rather have data pooled from across multiple exchanges and asset classes. Couple this with the fact that these systems can enter, exit, or cancel orders within fractions of a second, and it draws into question the fairness of the use of these systems in traditional markets.

---

*See source document for further definition.

^ We will assume that Voluntariness (individual autonomy) is a given, and so will not be addressing that prerequisite in the scope of this discussion.

As noted by the SEC (2010): "The speed of trading has increased to the point that the fastest traders now measure their latencies in microseconds. Is it necessary or economically feasible for long-term investors to expend resources on [..] these systems? If not, does the fact that professional traders likely always will be able to trade faster than long-term investors render these equity markets unfair for these investors?" An argument against this statement could be that for longer term-investors, entry *speed* is not so important, as a longer-term trading strategy is aiming to capture much larger market movements. As such, the difference between an entry speed of one millisecond and an entry speed of one second (or even one minute) is negligible. However, if we were to replace "longer-term investors" with retail investors/traders (who could be simply known as 'the little guys') the issue takes on a whole new form. In any case, it seems that the point of the SEC's remarks was that if a market participant cannot afford or otherwise gain access to these technologies, is it reasonable to expect that they could stand a chance of competing?

One answer to that question, is that it should be the responsibility of the trader to ensure that their own strategies are designed and tailored to suit their own circumstances. After all, suppose that a retail trader loses all of their money from voluntarily participating in the market. In this instance, it is self-evident that the removal of personal accountability alone is not grounds to justify setting the blame at the feet of High Frequency Traders.

Perhaps the strongest argument in favour of High Frequency Trading is that it increases liquidity in the markets in which it operates. This is due to the sheer volume of trades that can be executed within extremely short time frames. What this means for most ordinary participants is that there will almost always be orders available for them to enter or exit any small to medium sized position almost immediately, if they choose to do so. This helps the markets fulfil their fourth necessary function of Reliability. There is also a reasonable criticism of this argument in that this liquidity is not necessarily accessible to other market participants. This is what is known as 'ghost liquidity.' According to the Netherlands Authority for the Financial Markets (2016), ghost liquidity is "seeing liquidity, but then not being able to trade with it because it has suddenly disappeared."

Degryse et al. (2018) elaborate on this point, stating that "Ghost Liquidity exists when [high frequency] traders place duplicate limit orders on competing venues, intending for only one of the orders to execute, and when one does execute, duplicates are cancelled."

It is still debated how much of the total market liquidity is ghost liquidity, and whether or not it is a substantial enough amount to be of concern. However, it could be argued that just the mere existence of ghost liquidity could be detrimental to the health of the market. After all, what use is liquidity if participants cannot access it? In any case, while there is clearly a demonstrable benefit of HFT for retail traders in the form of liquidity (as a substantial amount of liquidity provided is actually accessible by other participants), the question still remains; "is that benefit overshadowed by the necessity of having to simultaneously 'compete' with these machines?"

2.1 "Dark Pools" and Hidden Orders

The second tenet of transparency stipulates that trade data should be shared in real-time so that all traders are on a level playing field and have the ability to respond accordingly.

Ordinarily, if a limit order is placed, that data is entered in to the books of the exchange and is visible to other market participants. And as the name suggests, hidden orders (not to be confused with 'iceberg orders') are not visible to other market participants until after the trade has been executed. Dark pools, on the other hand, are private financial forums that facilitate the exchange of large order block transactions that take place outside of the purview of the general public.

According to Buti et al., (2015) there are currently over 50 dark pools operating in the U.S., the 19 dark pools for which there is data available account for 14% of consolidated volume. It remains a point of contention as to whether or not the existence of these private pools of liquidity are fair to the users of regular

exchanges. In light of these facts the question then becomes, if a trading algorithm is allowed to run (that is, allowed to execute orders) without those orders being displayed to the rest of the participants, is that lack of transparency doing harm to the effectiveness of the system? One study on the effect of market transparency on market effectiveness (Madhavan et al., 2005) concludes in saying '[…] greater opaqueness may benefit markets that already offer a high degree of transparency.' Also noting that the opportunities to study these issues are quite rare and so will require further research.

3 INFORMATIONAL EFFICIENCY

The third tenet of Informational Efficiency states that the observed prices (of any given asset class) should incorporate all relevant information. Unfortunately, it does not specify that that information needs to be disseminated to all participants across the board in an overt manner, but only that it should be reflected in the price. From this, we can infer that it is the responsibility of the trader to research and become familiar with the markets in which they wish to participate. With that we can conclude that there is nothing 'wrong' with existence of hidden orders or dark pools or their use in the context of automated programs.

In the world of HFT, latency is extremely important. This subject is also closely related to Informational Efficiency, if we were to define it in terms of data transfer speeds (as opposed to the definition stated above).

The needs for high speed internet and even close physical proximity to the exchange have been key factors in reducing latency. These two factors would not be as readily accessible to the vast majority of retail traders, when compared with institutional or professional traders. However, while this may seem unfair to everyday traders, it may actually work in their favour. For a retail trader who only seeks to enter (or exit) relatively small positions at interspersed time frames, less aggressive liquidity providing HFTs could actually be quite useful.

In a study looking at the profitability of HFTs (Baron et al., 2012), they were divided in to three categories; aggressive (where >60% of trades are liquidity taking), mixed (where between 20% and 60% of trades are liquidity taking), and passive (where <20% of trades are liquidity taking). This study found the most aggressive strategies to be the most profitable by a significant margin. It was also found that newer entrants were not as profitable as existing firms, and had a higher probability to exit the market. This is likely due to the fact that established firms have faster systems and lower latency than their competitors, thus reducing their profitability and potentially forcing them out of the market. It is a point of concern that this will inevitably lead to a concentration of power in the field, and the value of these endeavours in relation to the overall health of the market is questionable.

For the less aggressive (and less profitable) HFTs who mostly provide liquidity instead of taking it, their contribution to the health of markets is far more obvious. Here, we can take the role of a Market Maker for example. This strategy (or service) seeks to provide a constant stream of liquidity to a host market with the aim of profiting from the bid-ask spread.

If XYZ stock is trading at $10 a share, the Market Maker will maintain a constant buy offer of $9.95 per share, and a sell offer of $10.05 per share. When a buy and sell order has been matched, the trades will execute and the Market Maker takes the difference in profit.

The service provided by Market Makers is particularly important for smaller stocks that may not otherwise have access to a large enough pool of liquidity. Thus, by exploiting this inefficiency Market Makers have found a way to make markets more efficient, albeit at a relatively small cost to the average trader.

Hasbrouck and Saar (2013) defined latency as 'the time it takes to learn about an event (e.g., a change in the bid), generate a response, and have the exchange act on the response.' Additionally, it was also noted that the fastest HFTs have reduced this entire process to as little as 'a couple of milliseconds'.

As mentioned previously, relative speeds of this magnitude are most desirable for (and therefore most likely employed by) those predatory, liquidity draining HFTs. A couple of questions that ought to be asked on this matter are; "Are aggressive HFTs harmful to the health

of their host markets? If so, what actions should be taken to ensure that they comply with the tenets of a healthy market?"

Roberts (1967) proposed that market efficiency should be broken in to three distinct levels depending on which information set is being considered. These are the three levels of the Efficient Market Hypothesis (EMH) as outlined by Malkiel (1989):

- Weak form EMH; prices fully reflect information derived from historical price action.
- Semi-strong form EMH; current prices not only reflect historical price information but also all publicly available information relevant to a company's securities.
- Strong form EMH; all information that is *known* to a market participant about a company is fully reflected in market prices.

While this hypothesis is well-defined, it has been criticized for not taking in to account things like behavioural psychology and its role in determining price action. Furthermore, at the time of its conception, high frequency trading as we know it today did not yet exist, and so also fails to consider the machinations of thousands of automated programs all with different aims and objectives.

Today, as our ability to collect and analyse data becomes ever more sophisticated, we may be able to leverage these technologies to monitor our markets at both micro and macro levels. This may provide us with the insights we need to more effectively monitor and regulate these markets going into the future.

4 CONCLUSION

Ultimately, it seems we may need to face the possibility that the debate surrounding the regulation and fairness of HFT currently has no straight forward answer. Trying to separate predatory and arguably less valuable HFTs from their passive counterparts would entail the use of yet another automated program. Whether or not banning aggressive HFTs would be desirable or warranted is another matter altogether.

In the words of Wellman and Rajan (2017), 'establishing laws and policies is a cumbersome and often slow process […] and may be expected to lag behind the pace of technological development'.

As such, it would seem reasonable to conclude that we would require an automated method of monitoring and regulating markets, as it is self-evident that our ability to regulate the markets through regular means could not possibly keep up with the rapid pace of innovation. However, this solution would undoubtedly give rise to even more questions and debate around the ethics of automation.

REFERENCES

James J. Angel and Douglas McCabe. 2013. Fairness in Financial Markets: The Case of High Frequency Trading. *Journal of Business Ethics.* https://doi.org/10.1007/s10551-012-1559-0.

Ricky Cooper, Michael Davis, Andrew Kumiega and Ben Van Vliet. 2020. Ethics for Automated Financial Markets. *Handbook on Ethics in Finance.* https://doi.org/10.1007/978-3-030-00001-1_18-1.

Frank J. Fabozzi, Sergio M. Focardi and Caroline Jonas. 2011. High-Frequency Trading: Methodologies and Market Impact. *Review of Future Markets.* https://www.theifm.org/sites/default/files/inline-files/High%20Frequency%20Trading-pub2011_0.pdf.

U.S. Securities and Exchange Commission. 2010. Part III: Concept release on equity market structure; Proposed Rule, 17 CFR Part 242. *Federal Register,* 17(13):3605.

Netherlands Authority for the Financial Markets (AFM). 2016. A Case Analysis of Critiques on High-Frequency Trading, page 7.

Hans Degryse, Rudy De Winne, Carole Gresse and Richard Payne. 2018. High frequency trading and ghost liquidity. *35th Annual Conference of the French Finance Association*. https://hal.archives-ouvertes.fr/hal-01894838.

Sabrina Buti, Barbara Rindi, and Ingrid M. Werner. 2015. Dark pool trading strategies, market equality and welfare. *Journal of Financial Economics,* 124:244-265.

Ananth Madhavan, David Porter and Daniel Weaver. 2005. Should securities markets be transparent?, *Journal of Financial Markets*, 8:266-288.

Matthew Baron, Jonathan Brogaard and Andrei Kirilenko. 2012. The Trading Profits of High Frequency Traders, pages 2-4.

Joel Hasbrouck and Gideon Saar. 2013. Low-Latency Trading. *Journal of Financial Markets,* 16:646-679.

Harry V. Roberts. 1967. Statistical versus clinical prediction of the stock market. *Unpublished manuscript,* University of Chicago.

Burton G. Malkiel. 1989. Efficient Market Hypothesis, pages 127-134.

Michael P. Wellman and Uday Rajan. 2017. Ethical Issues for Autonomous Trading Agents. *Minds & Machines*, 27:609-624. https://doi.org/10.1007/s11023-017-9419-4.