Algorithmic Trading: Balancing Automation and Regulation

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Nathaniel J. Gleberman

Spring 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Caitlin D. Wylie, Department of Engineering and Society

INTRODUCTION

On October 19th, 1987, global stock exchanges plummeted with the United States' market falling over 20% in mere hours (Dolan, 2023). Termed "Black Monday," this crash was caused by a number of economic factors alongside cascading stop-loss orders, algorithms that automatically sell positions at a designated percent loss (CFI Team, 2023). Market crashes like Black Monday can result in financial disasters such as increased unemployment, widespread poverty, and excess debt that can harm the lives of thousands (Investor.gov, n.d.). Algorithmic trading refers to computerized strategies for stock trading and encompasses programs such as the stop-loss orders (SEC, 2020). This paper seeks to understand the place of algorithms in the stock market with a focus on their potential flaws that might cause or enhance market crashes like Black Monday. By analyzing algorithmic trading, market crashes can be prevented while promoting a healthier marketplace that enhances the lives of investors.

To fully contextualize the place of algorithmic trading in the stock market, Actor-Network Theory (ANT) helps analyze the different relationships between human actors and nonhuman actors (Latour, 1992). In 2022, algorithmic trading was valued at 14.42 billion dollars and is expected to grow to 23.74 billion by 2027 (Mordor, 2022). This industry accounts for 60-73% of equity trading in the United States, meaning the vast majority of the stock market is now controlled directly by algorithms.Though nonhuman, algorithmic trading is developed by different human actors ranging from personal traders to companies of varying sizes. These human actors develop algorithmic trading strategies that prioritize profits by deciding when to buy and sell stock. All these strategies are connected as competitors inside the stock market, which serves as the network in ANT. This network is highly complex as the stock market is a collection of thousands of companies all with different missions, industries, and areas of influence. Money invested in the stock market is ultimately money invested in these companies and their practices, whether ethically right or wrong. This contextualization helps determine authority within the stock market, and these points of authority are where policy can positively impact the overall network of the stock market regardless of complexity.

Within the network, human actors can directly impact algorithmic trading and the companies listed on the stock exchange through policy and regulation. For example, in the aftermath of Black Monday, the U.S. Securities and Exchange Commission (SEC) set up circuit breakers, procedures for the automatic halting of trading given rapid drops in the market (Reissner, 2023). The SEC is the government authority responsible for ensuring fair markets for investors and companies alike making them a primary authority point in the network. In 2020, the SEC published a staff report on automated trading in U.S. capital markets to analyze benefits and risks of algorithmic trading. As an extensive report mandated by Congress, the SEC staff reached two conclusions: algorithmic trading promotes market health by increasing liquidity provision in normal conditions while some types of algorithms can enhance volatility in abnormal market conditions (SEC, 2020). The SEC focuses on monetary impacts and fails to fully evaluate effects on other important factors in the network, such as environmental impacts, ethical standards of companies, and even the mental health of developers.

ENVIRONMENTAL CONCERNS

Employing machine learning (ML) and artificial intelligence (AI), trading algorithms are frequently developed around data and powerful processes requiring large quantities of electricity to operate and maintain. The rise and prominence of AI platforms and large language models (LLMs) such as OpenAI's ChatGPT have brought the environmental concerns of AI to light. A study conducted at the University of Massachusetts found that training on a graphics processing unit (GPU) had the capability to generate nearly 625,000 CO_2e lbs of emissions depending on the ML model used (Strubell, Ganesh, & McCallum, 2019). A typical natural language processor (NLP) pipeline with tuning and experiments accounted for 78,468 CO_2e lbs. In comparison, the average American annually consumes around 36,000 CO_2e lbs. During model development, research, and tuning, these models can be repeatedly trained and rerun. Given this repetitive system, creation of these models is a highly energy consumptive process regardless of the baseline ML model used. These energy consumption concerns pass on to algorithmic trading through their use of ML, AI, and other automated processes running large scale data manipulation and predictive programs.

Electricity consumption is amplified further by the competitive landscape surrounding algorithmic trading. In an industry like quantitative stock trading where models of today can become ineffective tomorrow, constant research and model development are taking place. Adding to the high levels of energy consumption, algorithmic trading and artificial trading models are not limited to natural language processing alone. NLP could be a small part of the whole puzzle such as using ChatGPT as a sentiment scorer in a holistic model that looks at sentiment statements constantly scraped from the internet. The computational space and power needed is greatly expanded as strategies become more complex. Different hardwares exist, but more specialized and efficient hardwares could further save on resource consumption. Meanwhile, the ever present demand for human actors to find a competitive edge results in the vast electrical expenditure that compounds the energy consumption and impact on the environment by algorithmic trading.

Sacrificing the environment to trade stocks is a silly prioritization, so the necessity for better practices in the development and usage of algorithmic trading is paramount. Algorithmic trading is already a large consumer of energy in an industry constantly growing and competing daily. To lower emissions, algorithmic traders, especially businesses built around these algorithms, should be required to turn to renewable energy sources for part of their calculations. Through their control of the network, regulators like the SEC have the ability to reshape the impacts of algorithmic trading and companies listed on the stock market. This transition does not need to be 100%, but running all computations on nonrenewable energy sources is far more detrimental than running at least 5% on cleaner energy sources. The average consumer is likely familiar with a publicly traded company such as Apple, and being under the public eye can bring pressure to transition towards carbon neutrality. Meanwhile, quantitative funds such as Charlottesville's Quantitative Investment Management are not common household names and can avoid that same public scrutiny for better environmental practices. As such, I would encourage regulators to push for better environmental practices in the algorithmic trading industry. The SEC has made this necessity apparent in their March 2024 release of updated clauses for the environmental scoring of listed companies; however, the environmental concerns also fall on the companies allowed to trade stocks with algorithms.

This recent SEC addendum demonstrates that aside from the energy exploited by the hardware, sustainability concerns arise from the companies algorithms choose to invest in. Sustainability is the ability to keep a process going and for this paper, sustainability is directed towards the environment. A common ML model used by amateurs and professionals alike is the neural network, an ML model meant to simulate the neurons of a human brain. This advanced model takes as input basic quantitative stock variables such as price and volume. With inputted variables only pertaining to financial data, there is no way for an algorithm to measure the ethics of a company including their environmental impact. There is little incentive for human actors to design profit taking algorithms that consider scores not directly indicative of profits. As such, algorithms can freely determine that a company mining lithium with unpaid workers is as worthy an investment as a company that guarantees workers are fairly compensated. If unchecked by humans, algorithms have no programmed concern for humanity because algorithmic trading strategies are designed to make profits (Vellaiparambill & Natchimuthu, 2022).

Creating accountable algorithms is a difficult undertaking, so a better approach is to incentivize the developers shaping the algorithms. Researchers at Christ University in India propose including variables in the ML models that look at the environmental impact of companies such as Environmental and Social Governance (ESG) scores (Vellaiparambill & Natchimuthu, 2022). The SEC's environmental clauses only enhance this ability to measure environmental impact quantitatively company by company. Moreover, the Christ University researchers were able to show that doing so can be an effective profit making strategy comparable to the neural networks considering basic quantitative variables. This basic proof of profitability reveals that profit and sustainability are far from mutually exclusive. Companies that trade algorithmically should promote sustainability by prioritizing investment in companies that also prioritize sustainability.

Though a step in the right direction, socially responsible investing is a more nuanced issue than adding factors to trading algorithms. Not all algorithms use machine learning, so not all algorithms can turn to variables like ESG scores. Having humans confirm every trade would enhance the human connection to trading, yet some algorithms like high frequency traders (HFTs) work so fast humans physically cannot confirm every trade (Pavlus, 2019). To resolve this, I would propose required human screening for algorithmic traders prior to even permitting their algorithms to consider a company. Companies suspected of violating basic human rights or working against sustainability should not be served as inputs into algorithms in the first place. This way, HFTs and other algorithms only trade stocks that have been screened by humans. This data cleaning would maintain the humanity within algorithms while saving computational power through the elimination of certain companies and limiting of the initial data size to be processed.

OPACITY AND MARGINALIZATION

Regulating algorithmic trading is a challenging endeavor because the automated strategies are widespread, varied, and potentially unclear. Algorithmic trading has begun to adopt a more anthropomorphic approach by using ML models like deep neural networks to better identify trading opportunities. Dubbed artificial trading, these models represent a growing area of knowledge that researchers and developers admittedly do not fully understand. Deep neural networks come with opacity that clouds the ultimate reason behind each decision, and this opacity can perpetuate bias and elevate the difficulty of monitoring the ethics of artificial trading (Borch & Hee, 2022). Using technology that we do not fully understand is a concerning dilemma regardless of profession because the unknown leaves room for harm. This difficulty is not distinct to deep neural networks. Given the sheer quantity of trades provided by HFTs, layers of trades make analyzing market crashes difficult as demonstrated with the 2010 flash crash where some professionals believe further investigation is still needed to understand the crash (DeBold, n.d.). Opacity from algorithms makes regulation difficult because a proper analysis of the network becomes more challenging. This raises the question of whether or not human actors still control the network if the nonhuman actors can act in such a manner that years are needed to catch up with events in the stock market. Given this opacity, the algorithms themselves become an authority point for as long as the SEC fails to regulate them.

Looking at opacity, advocates of algorithmic trading can argue there is no difference from human trading; just as humans do not need to justify their trades, neither do these human mimicking algorithms. Many proponents of algorithmic trading believe the elimination of emotion is far more beneficial to the market than humans limited by greed or fear (Grimste, 2023). These arguments are correct in describing advantages of algorithmic trading, yet they fail to appreciate the sheer speed and volume with which these decisions are made. HFTs are an algorithmic trading strategy that utilizes the speed of a computer to ride the increase in price of large institutional trades that are broken into multiple timeframes ("High-Frequency Trading," n.d.). They often involve placing hundreds or thousands of trades in quick bursts to make money from discrepancies in the market. HFT's are the poster child for the algorithmic trading debate, often given a bad reputation for their seemingly exploitive strategy. Though controversial, data indicates HFTs promote a healthier marketplace by increasing liquidity and sacrificing monetary gains on actual stock positions for monetary gains on associated transaction rebates (Pavlus, 2019). Additionally, the SEC affirms algorithmic trading is a healthy market activity by increasing liquidity and helping satisfy orders (SEC, 2020). HFTs help reveal the marginalization of human actors within the stock market by unveiling public disapproval of the predatory nature of the algorithms in juxtaposition to professional approval on benefits to market liquidity.

HFTs and algorithmic trading promote a healthier stock market, but strong discrepancies exist in terms of accessibility. HFTs are an advanced strategy that involve connecting multiple systems and constantly analyzing market conditions for opportunities. Your average programmer is not running HFTs on their 2019 laptop. Though deemed "healthy" for their ability to promote market volume, more research needs to be conducted on who is losing in these situations because money made by HFTs does not appear from thin air. Companies capable of this technology are capable of benefitting from it, so smaller firms and retail investors frequently bear the price difference (Warren, 2013). The SEC is meant to create a fair market, and a fair market should entail equal accessibility to profits.

Like HFTs, algorithmic trading strategies can marginalize human traders through their rapid execution speed and knowledge barrier to entry. These walls are concerning because algorithms will always prevail over humans in reaction speed to good or bad news and quickly developing market trends. Many opponents of algorithmic trading argue that these systems have caused flash crashes, events where the stock market rapidly and suddenly plummets. Looking at the May 6th flash crash of 2010 again, academics and professionals alike found that the trace of the crash could not be properly analyzed because of the added complexity of algorithmic trading (Svetlova, 2022). Though the ability for algorithmic trading to perpetuate market sell-off remains unclear in historical crashes, the connection between the humans in the stock market and the nonhuman algorithms is abundantly present. Future market manipulation and flash crashes need to be closely monitored to better understand the power these nonhuman actors have in the stock market. The opacity needs to be understood to better mitigate human marginalization.

The barrier to entry for these algorithms is high, consisting of knowledge on the inner workings of both advanced real time programming and the stock market. In their paper *The Society of Algorithms*, Jenna Burrell and Marion Fourcade (2021) reveal a distinction between the human actors in the stock market. The coding elite can best be described as firms or individuals capable of creating computerized algorithms. As Burrell and Fourcade note, they are the ones “who hold and control the data and software” (2021, p. 215). The cybertariat is then the people producing and refining the data for the algorithms and run the risk of assisting in the automation of their own jobs. Everyone else is lumped into a large category of non-coders. With a heavy Marxist bias, *The Society of Algorithms* provides an abundance of sources defending their separate categorizations of human actors behind algorithms. I would take their grouping a step further for the stock market and distinguish non-coders as non-coders with authority over developers and non-coders without that authority. These distinctions between human actors will serve to highlight key differences dependent on knowledge and vested interest in the market.

The ability to develop algorithms to trade in the stock market creates a competitive leverage over non-coders because algorithms can deal with larger quantities of data. The sheer dominance of algorithmic trading in the stock market reveals its competitive edge. The coding elite are able to access and analyze the measurable parts of society that can directly correspond to the stock market's eventual movement. As most stocks are traded with algorithms, the coding elite who shape and understand these algorithms become the inadvertent dominant player in the market making them another authority within the network. Their authority is then captured by the non-coders in charge of them. Here is where the true marginalization takes place. Given these distinctions, authority over the market lies primarily with regulators and managers of the coding elite. The cybertariat and non-coders without power appear entirely powerless in comparison to these authority points. A responsibility to protect these different groups then lies with both the regulators and managers. Lastly, the coding elite have authority in their ability to create the algorithms and imbed them with authority; however, as the next section seeks to highlight, the coding elite are subject to the regulators and non-coding managers.

MENTAL STRAIN

The coding elite need their product to work effectively, yet competition might prevent this. Trading strategies are constantly being developed and reworked when deemed no longer effective. In a joint paper from MIT and Northwestern, S.P. Kothari and Robert Pozen (2023) elaborate on this dilemma which they refer to as a paradox. Given the impermanence of new strategies and the growing ability of machine learning and artificial intelligence to replicate results, constantly developing strategies is a costly and unsustainable approach to distinguishing competitors. Writing for a business journal, Kothari and Pozen conclude that the sustainable approach to thriving in the realm of algorithmic and artificial trading is ultimately customer service. I would agree and find that no matter how non-human the inner working of the stock market becomes, this network is still a means of tracking human nature through the economy and the humans behind the algorithms and investments. There is an almost ironic inability for these algorithms to escape humanity that should serve to emphasize the need to protect the humans behind their development.

This constant paradox of competition that demands the reshuffling of strategies and staff creates a lot of developmental strain on the coding elite and the cybertariat as working hours escalate and stress levels elevate. In an internal report from Goldman Sachs, employees were averaging 5 hours of sleep per night with 98 hours of work per week (Clark, 2023). The finance industry pays heavily for this ability to work relentlessly, but this is an unsustainable lifestyle where the coding elite seemingly become victim to the demand for their coding skills. Not only should algorithms be monitored for their power in the market, the firms should be monitored by the SEC for appropriate work place conditions and the collection and analysis of data regarding the coding elite. Companies should practice and enforce safe approaches to competition that do not isolate anyone in the development process from the data collectors to the top-level programmers. The algorithmic trading industry needs to be protected both externally and internally to promote the health of the stock market and the health of the human actors behind the algorithms.

CONCLUSION

Excessive resource consumption, disparities between human and machine reaction times, and heavy mental strain on developers are all concerns of algorithmic trading to monitor. Increased volume and market activity are major proven benefits, yet added opacity makes attributing negative impacts difficult. Algorithmic trading is rapidly changing especially in the face of huge growth in the AI sector and the market itself. More predatory strategies could be developed or hidden better, and the SEC tends to be slow to identify and react to problems caused by algorithms in the market. Cooperation between actors with authority such as the coding elite, non-coding managers, the algorithms themselves, and regulators is necessary to create a stock market that is fair and accessible for all.

REFERENCES

Borch, C., & Hee Min, B. (2022). Toward a sociology of machine learning explainability: Human–machine interaction in deep neural network-based automated trading. *Big Data & Society*, *9*(2), 205395172211113. https://doi.org/10.1177/20539517221111361

Burrell, J., & Forcade, M. (2021). Society of Algorithms. *Annual Reviews*, *47*, 213–237. https://doi.org/10.1146/annurev-soc-090820-020800

CFI Team. (2023, October 18). *Black Monday Market Crash*. Corporate Finance Institute. https://corporatefinanceinstitute.com/resources/equities/black-monday/

Clark, Ken (2023). Making It Big on Wallstreet. *Investopedia*. https://www.investopedia.com/articles/financial-careers/09/big-career-on-wall-street.asp#toc-how-stressful-is-it-to-work-on-wall-street

DeBold, R. K., Bradley Hope and Tynan. (n.d.). *"Flash Crash" a Perfect Storm for Markets*. WSJ. Retrieved May 5, 2024, from https://graphics.wsj.com/flash-crash-timeline/

Dolan, B. (2023, April 30). *What caused Black Monday, the 1987 stock market crash?*. Investopedia. https://www.investopedia.com/ask/answers/042115/what-caused-black-monday-stock-market-crash-1987.asp

Grimste, C. (2023, September 25). *Algorithmic trading - overview, Examples, pros and cons*. Wall Street Oasis. https://www.wallstreetoasis.com/resources/skills/trading-investing/algorithmic-trading

High-Frequency Trading. (2019, October 29). High-Frequency Trading. *CFA Institute*. https://rpc.cfainstitute.org/en/policy/positions/high-frequency-trading

Latour, B. (1992). Where are the missing masses? The sociology of a few mundane artifacts. *Shaping technology/building society: Studies in sociotechnical change*, *1*, 225-258.

Lezginov, M. (2022, January 4). *AI Trading vs. Algorithmic Trading. What's the Difference?* Scopic. https://scopicsoftware.com/blog/the-future-of-ai-trading/

Mordor Intelligence. (2022). Algorithmic Trading Industry Report. *Mordor Intelligence*. https://www.mordorintelligence.com/industry-reports/algorithmic-trading-market

Pavlus, J. (2019). Do high-frequency traders deserve their bad rap? *Northwestern University*. https://insight.kellogg.northwestern.edu/article/do-high-frequency-traders-deserve-their-bad-rap

Pozen, R. C., & Kothari, S. P. (2023). The Paradox of AI and Investing. *International Banker*. https://internationalbanker.com/technology/the-paradox-of-ai-and-investing/

Reisner, R. (2023, December 27). *8 of the biggest stock market crashes in history - and how they changed our financial lives*. Business Insider. https://www.businessinsider.com/personal-finance/biggest-stock-market-crashes-in-history?op=1

Securities and Exchange Commission (2020), Staff Report on Algorithmic Trading in U.S. Capital Markets 1–92. https://www.sec.gov/files/Algo_Trading_Report_2020.pdf

Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (pp. 3645–3650). Florence, Italy: Association for Computational Linguistics.

Stock Market Circuit Breakers. Investor.gov. (n.d.). https://www.investor.gov/introduction-investing/investing-basics/glossary/stock-market-circuit-breakers

Svetlova, E. (2022). Ai Ethics and Systemic Risks in Finance. *AI and Ethics*, *2*(4), 713–725. https://doi.org/10.1007/s43681-021-00129-1

Vellaiparambill, A., & Natchimuthu, N. (2022). Ethical tenets of stock price prediction using Machine Learning Techniques: A sustainable approach. *ECS Transactions*, *107*(1), 137–149. https://doi.org/10.1149/10701.0137ecst

Warren, Todd K., (2013, January 22). High Frequency Trading – Issues Facing Investment Firms. *Counsel, Sadis & Goldberg LLP.* https://www.marcumllp.com/insights/high-frequency-trading-issues-facing-investment-firms