THE CAPCO INSTITUTE

JOURNAL

OF FINANCIAL TRANSFORMATION

TECHNOLOGY

Understanding the implications of advanced AI on financial markets

MICHAEL P. WELLMAN

GenAI

60TH EDITION

2024/2025 EDITION

a wipro company

THE CAPCO INSTITUTE

JOURNAL OF FINANCIAL TRANSFORMATION

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

CONTENTS

TECHNOLOGY

ORGANIZATION

REGULATION

CAPCO CEO WELCOME

DEAR READER,

Welcome to our very special 60th edition of the Capco Journal of Financial Transformation.

The release of this milestone edition, focused on GenAI, reinforces Capco's enduring role in leading conversations at the cutting edge of innovation, and driving the trends shaping the financial services sector.

There is no doubt that GenAI is revolutionizing industries and rapidly accelerating innovation, with the potential to fundamentally reshape how we identify and capitalize on opportunities for transformation.

At Capco, we are embracing an AI infused future today, leveraging the power of GenAI to increase efficiency, innovation and speed to market while ensuring that this technology is used in a pragmatic, secure, and responsible way.

In this edition of the Capco Journal, we are excited to share the expert insights of distinguished contributors across academia and the financial services industry, in addition to drawing on the practical experiences from Capco's industry, consulting, and technology SMEs.

The authors in this edition offer fresh perspectives on the mindful use of GenAI and the implications of advanced GenAI on financial markets, in addition to providing practical and safe frameworks for boards and firms on how to approach GenAI governance.

The latest advancements in this rapidly evolving space demonstrate that the potential of GenAI goes beyond automating and augmenting tasks, to truly helping organizations redefine their business models, processes and workforce strategies. To unlock these benefits of GenAI, I believe that firms need a culture that encourages responsible experimentation and continuous learning across their organization, while assessing the impact of the potential benefits against a strategic approach and GenAI framework.

I am proud that Capco today remains committed to our culture of entrepreneurialism and innovation, harnessed in the foundation of our domain expertise across our global teams. I am proud that we remain committed to our mission to actively push boundaries, championing the ideas that are shaping the future of our industry, and making a genuine difference for our clients and customers – all while ensuring to lead with a strategy that puts sustained growth, integrity and security at the forefront of what we do.

I hope you'll find the articles in this edition both thought-provoking and valuable as you create your organization's GenAI strategy and future direction. As we navigate this journey together, now is the time to be bold, think big, and explore the possibilities.

My greatest thanks and appreciation to our contributors, readers, clients, and teams.

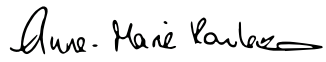

Annie Rowland, **Capco CEO**

UNDERSTANDING THE IMPLICATIONS OF ADVANCED AI ON FINANCIAL MARKETS

MICHAEL P. WELLMAN | Lynn A. Conway Collegiate Professor of Computer Science and Engineering University of Michigan, Ann Arbor

ABSTRACT

The rapid advancement of surprisingly capable AI is raising questions about AI's impact on virtually all aspects of our economy and society. The nexus of AI and finance is especially salient, building on the impact AI has already had on trading and other financial domains. New AI developments could exacerbate market manipulation, and otherwise create loopholes in regulatory regimes. Anticipating these potential impacts suggests directions for market design and policy that makes financial markets robust to advanced AI capabilities.

1. INTRODUCTION

It seems that everyone is in an excited state these days about the apparently rapid advances in artificial intelligence (AI), and its potential to solve big problems or create new ones. This excitement is warranted, on both sides. AI promises to bring us extraordinary benefits through new capabilities to expand knowledge and automate difficult tasks, and by making a variety of valuable services accessible and affordable to broad segments of our society. AI also threatens us with an array of potentially negative consequences, including risks to security posed by malicious exploitation of AI, risks to safety due to inadvertent AI behaviors, and the risk of systemic disruption to the ways we work and live. The promises and threats of AI pervade essentially every area of our economy and society, including quite distinctly the financial sector.

In this article, I focus on the nexus of AI and finance, and particularly on implications of advanced AI for financial markets. I describe at a high level how AI is employed in markets today, and some of the possible implications of the newest AI developments. Following brief background on algorithmic trading, I focus on three ways in which the latest AI technology may bring some new considerations for security, efficiency, and fairness of our capital markets.

Let us start with the necessary qualifier that the future path of advanced AI is highly uncertain. If somebody tells you they know where AI technology will be in five years or ten years – or even next year – be very skeptical. Technical breakthroughs are inherently unpredictable, and AI has a particular capacity to surprise. It has surprised us many times, most recently in Fall 2022 by ChatGPT and its ilk. Even experts with the deepest understanding of generative AI (GenAI) techniques such as large language models (LLMs) were surprised at the quality and utility of results they are able to produce. We are also sometimes surprised by limitations and weaknesses of AI technology, or roadblocks to advancement. Either way, AI is likely to keep surprising us.

Please also keep in mind that we have limited visibility into developments that are already in the pipeline. There are likely thousands of active projects aiming to harness the latest GenAI advances in novel products and services. Startup companies, corporate development teams, and public and private research labs around the world are all exploring how to put GenAI to work. Many of these will fail (or already have) but some are likely to surprise us with new capabilities and impactful use cases.

2. ALGORITHMIC TRADING

Under-the-radar development is actually the main story of deployment of AI in financial markets up to now. AI is already widely adopted in support of trading in markets, where it has had a significant impact. The shift to electronic markets over the past few decades has had many effects, notably on speed of reaction to information. One effect has been to enable implementation of algorithmic strategies developed using AI technology such as advanced machine learning. Whereas the term "algorithmic trading" does not necessarily entail that there is "AI inside", it is surely the case that developers of trading algorithms often employ cutting-edge AI techniques. I would even go as far as to claim that algorithmic trading represents the first widespread use of autonomous agents (i.e., AI decision making without humans in the loop) in a high-stakes and economically significant domain.

Electronic markets are well suited for software agents in part due to their simple and circumscribed interfaces (data feeds and order submission through well-specified protocols), which narrows the scope of agent behaviors that must be considered. Nevertheless, it may seem surprising that financial trading would be a first domain for autonomous operation, given the stakes involved, and thus the risks. It turned out that the advantages frequently outweigh the risks. Markets place a premium on the ability to process a multiplicity of information sources at high velocity, combined with rapid response time, both of which are in the wheelhouse for algorithms. In a situation where the first to respond to information captures the profit, putting a human in the loop is simply not an option. The returns to effective strategies are such that the research and development to produce them was worth the try, and once some initial success was demonstrated, regular processes and business models could be built around them.

Gauging the exact extent and nature of AI employed in algorithmic trading today is not possible, due to a lack of public information. Trading firms do not publish information about their strategies, for obvious proprietary reasons, and they also tend to be extremely protective about information regarding broad approaches, technology employed, data and information sources, and really everything about their strategic methodology and operations. Nevertheless, there are exceptions, and some information occasionally leaks out or is inferable from hiring practices, technology investments, or market observations. As a result, we can be quite confident about the general assessment that use of cutting-edge AI for trading is pervasive in current financial markets.

The opacity of state-of-the-art trading technology is itself one source of risk. There exists a keen public interest in understanding how various trading practices affect the fairness, efficiency, and stability of financial markets. The need for open information on AI trading strategy was a major motivation for my own group's research in this area. I should emphasize that the goal of this research – by us or others – is not to assess whether algorithmic trading in general is beneficial or harmful to financial markets. The goal of the research is to tease apart the practices and circumstances that help or hurt, and further to identify market designs or regulations that promote the beneficial practices and deter the harmful ones.

For example, we have found that algorithmic market making improves efficiency and can be beneficial to those trading for investment, particularly when markets are thin and the market makers are competitive [Wah et al. (2017)]. In thick markets, though, algorithmic market making can extract surplus from investors. Another issue that we have investigated is "latency arbitrage": the deployment of practices that leverage miniscule advantages in response time, measured in milliseconds or microseconds, to extract profit from trades that would have happened anyway [Wah and Wellman (2016)]. We and others have advocated for a mechanism called frequent batch auctions, where markets clear at fixed intervals, such as every half-second, rather than continuously, to short-circuit the latency arms race, thus improving both fairness and efficiency [Budish et al. (2015)].

3. THE NEWEST AI

While there is still much we need to understand about today's algorithmic trading and its effects, the latest AI developments are raising qualitatively new issues about the implications for financial markets. The pace of technical advance in AI has been quite astounding in the last decade or so, but at the risk of over-simplification let me focus on two broad categories.

- **Deep reinforcement learning (DRL):** the use of deep neural networks to represent strategies, trained using reinforcement learning [Sutton and Barto (2018)]. This approach has demonstrated enormous advances over the past decade. DRL was the technology behind Google DeepMind's breakthroughs in the game of Go [Silver et al. (2016)] and protein folding [Jumper et al. (2021)] (recently recognized with a Nobel Prize), for example, and indeed was the basis for DeepMind's original formation. DRL is particularly salient for algorithmic trading because

it enables the partial or full automation of strategy generation. That is, with DRL one can train a strategy that responds to market information with actions without any human expressly programming the logic of this response.

- **Large language models (LLMs):** the massive neural networks trained to generate fluent natural language responses to textual prompts [Zhao et al. (2023)]. LLMs are the technology behind chatbots (e.g., OpenAI's ChatGPT), and part of the broader category of GenAI methods that have sparked the current explosion of interest in AI. LLMs are especially significant because they open up the language channel. That is, they enable AI methods to interact with humans (or other AI systems) in natural language, as well as standard computer languages. This allows them to be deployed in situations with open-ended interfaces, not just environments explicitly crafted for programmatic interaction.

These two categories are not entirely separate. In fact, configuring LLMs to perform useful tasks requires shaping how they respond to prompts using DRL, specifically training with reward signals based on human feedback. Combining the power of massive pre-trained models with DRL is indeed one of the most promising approaches for the next generation of autonomous agents.

At a high level, both of these new AI capabilities carry the potential to dramatically extend the autonomy and scope of algorithmic trading. Automating the strategy generation process itself adds a level of autonomy, in the sense of shifting the human control to a more indirect and abstract layer of supervision. Opening the language channel enables the trading agent to act autonomously in a much broader scope of situations. In principle, a capable chatbot could trade flexibly with human securities dealers in an over-the-counter trading environment.

4. MARKET MANIPULATION

Practices that inject misleading information about market conditions can seriously compromise the transparency and thereby the fairness and efficiency of public markets. Market manipulation is an old practice, but AI may turn out to amplify the power of would-be manipulators to at achieve their manipulative purpose, with lower cost and risk of detection. In response, sophisticated machine learning techniques can also be used by market regulators for enhanced surveillance, detection, and enforcement. Measures based on machine learning, however, are subject to countermeasures that aim to undermine or circumvent the learning system [Papernot et al. (2018)]. Use of machine learning for regulation naturally sets up what is called an adversarial learning situation [Vorobeychik and Kantarcioglu (2018)], a kind of AI arms race, between the detector and evader. An inherent feature of adversarial learning is that any advance in detection technology can be immediately exploited by the evader to improve its evasion. Where this leads in any given situation is an open question. In our market manipulation studies, we have found that evading detection also weakens the manipulation [Wang and Wellman (2020)], but whether that will always be the case we cannot be sure.

The first-order concern is that malicious parties could use AI intentionally to manipulate markets. It is also possible that AI-developed trading algorithms could produce strategies that employ manipulation or other harmful tactics, even if such manipulation was not the specified objective. In fact, our research has demonstrated the possibility of an AI independently learning to manipulate a financial benchmark, given only the objective of seeking profit [Shearer et al. (2023)]. Are current regulations regarding market manipulation adequate to handle such a situation? Much of the existing law depends on "intent" to manipulate, and how that would apply to an AI algorithm that learned manipulation on its own is unclear.

This is just one example of what I call an "AI loophole". Our existing laws, generally speaking, are written based on the assumption that it is people who make decisions. When AIs are deciders, do our laws adequately ensure accountability for those putting the AIs to work?

The second issue is specific to the advances in language processing exhibited by LLM-based systems like ChatGPT. Arguably, one of the reasons that AI has been so successful in financial trading already is that the interface to markets (streams of buy and sell orders) is so simple and circumscribed. Text processing techniques based on machine learning have also been employed in trading to some extent, but the new LLMs can potentially take this to a new level. With broad language competency, massive bodies of human-generated information become available as material that can be traded on.

The new models also provide the capacity to generate text, thus opening up new language channels for AI influence. With generative capacity, systems can actively query humans to elicit information that may not have been available otherwise.

They can also use this channel to inject misleading information, which brings us back to market manipulation. Just as human manipulators employ social media in their "pump-and-dump" schemes, we should expect efforts to amplify such messages using AI.

This manipulation concern is just a special case of the broader problem of misinformation and fraud. In the wrong hands, AI can be great technology for deploying scams. Of course, this issue is relevant well beyond the financial domain.

5. CONCENTRATION OF INFORMATION

The final issue I would like to raise here relates to how the new AI technology obtains its power through training over massive datasets. It appears that qualitative leaps in capability can come from large scale source information. A corollary is that only entities with access to such large bodies of information can produce AI systems with the greatest performance. In the realm of financial trading, this could mean that concentrations of information access and ownership could convey extraordinary advantages. This naturally raises questions about how trading on information aggregated at massive scale could affect fairness and efficiency of our financial markets.

There now exist great stores of non-public information about people and their activities that have been amassed by companies through provision of information services and other online interactions. Much of this is willfully provided to enable or improve the quality of services, and often with understanding that it may also be used for marketing or related purposes. Considerations about this are typically framed in terms of personal privacy. Protection of personal privacy is indeed an important concern, but it may be equally important to consider the strategic implications of aggregations of information, as they affect us collectively beyond the individuals associated.

6. CONCLUSION AND POLICY RECOMMENDATIONS

I have highlighted three broad issues in this article: AI loopholes, opening the language channel, and concentration of information. Each may affect the balance of power in markets, through development of "super-manipulation" capabilities, strategic advantage, or other means. These issues are just a few of the ways that new AI technologies pose novel concerns for financial markets. AI also offers the potential to protect market integrity and level the investment playing field. Which

> *"Our existing laws, generally speaking, are written based on the assumption that it is people who make decisions. When AIs are deciders, do our laws adequately ensure accountability for those putting the AIs to work?"*

effects predominate will be in large part determined by how we reconsider market designs and governance mechanisms for the world of AI-powered trading. I conclude with some recommendations about some tasks we should focus on in this endeavor.

Design markets for algorithms first. In the not-distant future, it is reasonable to expect that any interaction between a person (or organization) and a market will be mediated by algorithms. Even retail investors will have access to agents that can implement execution strategies on their behalf, rather than submitting limit or market orders directly. It is thus quite plausible to design market mechanisms under the assumption that the participants are algorithms. Though the algorithms themselves may have different computational resources and information sources, they will all have access to up-to-date market state and rapid response capabilities. Accordingly, it is possible to consider principles of fairness in availability to the interface, and fine-grained control of the extent and timing of information provided.

Conduct an inventory of laws and regulations to check for AI loopholes. To the extent that conduct proscribed for human actors could evade sanction through automation by algorithms, we have a potential AI loophole. It behooves financial regulators to consider where such situations might arise (e.g., for regulations expressed in terms of intent), and patch up the legal framework to prevent such evasion. Consider what new requirements may be necessary to ensure that behavior by algorithms is traceable to accountable parties, and areas of potential misconduct that are made newly practical thanks to AI.

Build foundations for trusted information. A functional financial system depends on widespread availability of reliable information, about markets, companies, assets, and the economy in general. Misinformation, including disinformation generated and promulgated through AI, pollutes the information environment and undermines sound financial decision making. AI itself will likely not be sufficient to counter misinformation, and so we would be wise to invest in mechanisms and associated infrastructure that could positively establish foundations of trust for critical financial information.

Support development of third-party evaluation tools. Systems developed using AI methods are generally more prone to unpredictable behavior, due to their complexity and their mode of development. In particular, algorithms using models trained via machine learning may behave in surprising ways in situations not covered by their training data. Subjecting them to rigorous testing and evaluation may increase confidence in their safety and performance. Having third parties develop the testing regimes reduces concerns about conflicts of interest, and the risk of blind spots due to coupling of design and evaluation. There are signs that the marketplace is starting to develop AI evaluation services, but this could be accelerated by development of standards and certification requirements.

Research. That an academic researcher would call for more research is perhaps the least surprising recommendation. But the truth is, there is a lot we have yet to understand about the implications of advanced AI on financial markets, especially the scope of the risks and effectiveness of preventive strategies. Creating the knowledge necessary to prepare for a financial system with AI is a compelling public interest.

REFERENCES

Budish, E., P. Cramton, and J. Shim, 2015, "The high-frequency trading arms race: frequent batch auctions as a market design response," Quarterly Journal of Economics 130:4, 1547–1621

Jumper, J., et al. 2021, "Highly accurate protein structure prediction with AlphaFold," Nature 596, 583–589

Papernot, N., P. McDaniel, A. Sinha, and M. P. Wellman, 2018, "SoK: Security and privacy in machine learning," 3rd IEEE European Symposium on Security and Privacy

Shearer, S., G. Rauterberg, and M. P. Wellman, 2023, "Learning to manipulate a financial benchmark," 4th International Conference on Artificial Intelligence in Finance, 592–600

Silver, D., et al., 2016, "Mastering the game of Go with deep neural networks and tree search," Nature 529:7587, 484–489

Sutton, R. S., and A. G. Barto, 2018, Reinforcement learning: an introduction, second edition, MIT Press

Vorobeychik Y., and M. Kantarcioglu, 2018, Adversarial machine learning, Morgan & Claypool Publishers

Wah, E., and M. P. Wellman, 2016, "Latency arbitrage in fragmented markets: a strategic agent-based analysis," Algorithmic Finance 5, 69–93

Wah, E., M. Wright, and M. P. Wellman, 2017, "Welfare effects of market making in continuous double auctions," Journal of Artificial Intelligence Research 59, 613–650

Wang, X., and M. P. Wellman, 2020, "Market manipulation: an adversarial learning framework for detection and evasion," 29th International Joint Conference on Artificial Intelligence, 4626–4632

Zhao, W. X., et al., 2023, "A survey on large language models," https://tinyurl.com/2eb54bbs

ABOUT CAPCO

Capco, a Wipro company, is a global management and technology consultancy specializing in driving transformation in the energy and financial services industries. Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on LinkedIn, Instagram, Facebook, and YouTube.

WORLDWIDE OFFICES

APAC
Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurugram
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST
Dubai

EUROPE
Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA
Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto

SOUTH AMERICA
São Paulo

THIS UNIQUE IMAGE WAS GENERATED USING MID-JOURNEY, STABLE DIFFUSION AND ADOBE FIREFLY

WWW.CAPCO.COM

CAPCO
a wipro company