Written Testimony before the US Senate Committee on Banking, Housing, and Urban Affairs

Hearing on “Artificial Intelligence in Financial Services”

Michael P. Wellman

Lynn A. Conway Collegiate Professor of Computer Science & Engineering
Richard H. Orenstein Division Chair of Computer Science & Engineering

University of Michigan
Ann Arbor, MI

20 September 2023

Chair Brown, Ranking Member Scott, and distinguished members of the committee, it is a privilege to testify before you today. The committee's attention to artificial intelligence and its implications for the financial system in particular is timely and important.

It seems that everyone is in an excited state these days about the apparently rapid advances in artificial intelligence, 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 my testimony this morning, I will focus on the nexus of AI and Finance, and particularly on implications of advanced AI for financial markets. I will aim to describe at a high level how AI is employed in markets today, and convey a general sense of the possible implications of the newest AI developments. Following some background on algorithmic trading, I will focus on three areas where the latest AI technology may bring some new considerations for security, efficiency, and fairness of our capital markets.

Before getting into the substance, let me provide some background on myself for context on where I am coming from. As noted in the introduction, I am Professor and Chair of Computer Science & Engineering at the University of Michigan. I earned my PhD in Artificial Intelligence from MIT in 1988. I have worked in the field as an AI researcher ever since, first in the US Air Force and for the last 30+ years on the faculty at the University of Michigan. I am known for research at the intersection of AI and economics, including contributions to the field of autonomous agents and multiagent systems, and applications to electronic commerce. I am a Fellow of the Association for the Advancement of Artificial

Intelligence (AAAI) and of the Association for Computing Machinery (ACM). For the past dozen years or so, my research has focused on understanding the implications of AI for financial markets and the financial system. Regarding this topic, I have served on advisory committees for the US Treasury Office of Financial Research (OFR), and the Financial Industry Regulatory Authority (FINRA). All opinions expressed in this testimony are my own, and not attributable to organizations I am or have been employed by, received funding from, or provided advice to.

I would also like to preface my remarks 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—don't believe them. Technical breakthroughs are inherently unpredictable, and AI has a particular capacity to surprise. It has surprised us many times, including within the last year by ChatGPT and its ilk. Even experts with the deepest understanding of generative AI techniques such as large language models (LLMs) were surprised at the quality and utility of results they can 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 generative AI 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 generative AI to work. Many of these will fail but some are likely to surprise us with new capabilities and impactful use cases.

Under-the-radar development is actually the story of the deployment of AI in financial markets up to now. In fact, 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.

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 high-level 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. 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. 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.

Let me now move more specifically to some newer issues posed by the latest AI developments.

The first issue is market manipulation. Practices that inject misleading information about market conditions can seriously compromise the transparency and thereby the fairness and efficiency of public markets. Of course market manipulation is an old practice, but AI can be more effective at achieving its 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. This naturally sets up what is called an "adversarial learning" situation, 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, but whether that will always be the case we cannot be sure.

What I have been discussing so far is the concern 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. 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. 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. This opens up massive bodies of human-generated information 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.

The final issue I would like to raise today 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.

The three issues I have highlighted here are 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 effects predominate will be in large part determined by how we reconsider market designs and governance mechanisms for the world of AI-powered trading.

This concludes my prepared remarks. I am grateful for the opportunity to present this perspective to the committee and welcome your questions.