Examining Misbehavior of Artificial Intelligence Traders in Simulated Financial Markets

Rahul Dasgupta, Class of 2024

When it comes to the implementation of Artificial Intelligence into financial markets, there are concerns that AI will learn to "misbehave" or develop and perform trading strategies that are ethically questionable or even illegal. Professor David Byrd showed that these concerns are not baseless in his study that revealed that, when given a simple profit maximizing goal, AI trading agents will learn the financial crime of *spoofing* the market it trades in (Byrd 2022).

In my research, I explored other scenarios in which AI agents may learn to misbehave or commit financial crimes. In particular, I focused on what misbehavior may arise with the use of AI in brokerage firms. Brokerage firms trade on behalf of clients that don't interact directly with an exchange, and they often make their own trades as well. Frontrunning is a financial crime committed by traders, such as brokers, when they place orders based on their knowledge of client orders that they are yet to execute. For my research, I aimed to examine whether brokerage firm AI trading agents may inadvertently commit frontrunning or frontrunning-like behavior.

Professor Byrd built a discrete event simulation program called ABIDES (Agent-Based Interactive Discrete Event Simulation) that has the capability of simulating a financial market, and I built upon his simulation in order to conduct my experiments. I created several types of agents having to do with brokerage firms, including "Broker agents" which handled orders, broker client agents which traded through Broker agents rather than directly with an exchange, and broker trading agents which traded on behalf of a brokerage firm. We assumed that due to order internalization (the process by which brokers match different client orders to each other in order to avoid transaction costs), that brokerage firms would maintain a list of client orders that are yet to be executed. We also assumed that a trader for a brokerage firm would have access to such a list in order to execute their own trades without transaction costs, and we believed that visibility of this list may lead AI trading agents within brokerage firms to commit misbehavior having to do with frontrunning.

We developed an AI broker trading agent which used deep reinforcement learning built with the TensorFlow and Keras Python libraries. The experiments I conducted had two versions of the AI agent: one with access to the internalized order list and one without. The first experiment included broker clients which placed orders entirely at random, as we believed that such orders would be representative of retail traders that make up a broker's clientele. The results of this experiment did not lead us to believe that the AI agent was committing frontrunning-like behavior as the trained agent would most often yield negative results, losing significantly on average, and this is most likely due to the randomness of the client trades. With a trading agent that had largely negative results, we deemed it unlikely that the agent was exhibiting exploitative behavior, and revised our approach to attempt to more accurately represent a brokerage firm's clientele.

The second experiment had broker clients that traded based on a spread of values for a believed "fundamental value" of what a stock's price should be worth, which we believed would be closer to retail trader behavior. In this experiment, the AI trading agent yielded much more positive results, however the general trends of the results of the agent with access to the internalized order list and the results of the agent without access were indistinguishable. However, close examination of the executed trades in a simulation day revealed that there are several cases where the AI agent with order list visibility places an order before a client's order of the same direction is sent to the exchange. There is a chance that this frontrunning-like behavior is coincidental and not learned, but it still constitutes a situation where a broker places an order before placing a client order in the same direction while having access to client order information. As brokerage firms incorporate AI into their trading strategies, these cases must be monitored and examined. Regulations for how an AI trader prepares orders may need to be implemented as this behavior may be a sign of exploitation by AI which might otherwise go undetected.

Faculty Mentor: Professor David Byrd

Funded by a J.P. Morgan AI Research Faculty Fellowship

Byrd, David. "Learning not to spoof." Proceedings of the Third ACM International Conference on AI in Finance, 2022, https://doi.org/10.1145/3533271.3561767.