A Work Project, presented as part of the requirements for the Award of a Master's degree in Finance from the Nova School of Business and Economics.

THE IMPACT OF ARTIFICIAL INTELLIGENCE ON THE DEMOCRATIZATION OF PROFITS IN FINANCIAL MARKETS

VASCO LEMOS BESSA PELETEIRO

53370

Work project carried out under the supervision of:

Pedro Souto

20/12/2023

Abstract

Through an agent-based simulation of the cryptocurrency market, the study assesses the ramifications of an AI-powered entity making predictive market interventions on the retail investor's ability to generate profits. Emphasizing simplicity and realism, the model reflects realistic market mechanics and diverse investor behaviours. The results reveal stark contrasts in returns, where regular investors experience negative returns post-AI agent introduction, while the AI agent consistently achieves significant profits. This study highlights potential challenges for individual investors in markets increasingly dominated by AI-equipped entities, emphasizing the need for discussions on fairness, transparency, and accessibility in finance amidst AI proliferation.

Keywords: Agent-Based Models, Financial Markets, Market Dynamics, Artificial Intelligence

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209).

Table of Contents

Table of Figures

Introduction

For those deeply immersed in finance, in the past two years a pertinent query often surfaced: how might AI influence financial markets? While AI-driven tools analysing market trends and offering passive investment guidance contribute positively to society, a different narrative emerges concerning a distinct type of agent.

In this period, the landscape of artificial intelligence has undergone a remarkable evolution, notably propelled by significant advancements in Large Language Models. This surge has led to the proliferation of thousands of new AI tools, signalling an ongoing trend toward further development in this domain. Either built in-house or developed by an AI lab, the development of agents capable of leveraging vast live data for highly precise active investment strategies, potentially even suggesting market manipulation tactics, will most likely be exclusive to the most powerful market figures.

The implementation of this very powerful type of AI agent within the tools that these players possess to optimize their financial performance is almost a certainty. It might not happen in the next 5 years but will surely happen in the next few decades. This reality then begs the question: what impact will this scenario have on the chances of individual investors collecting profits from financial markets? Said in another way: what is the impact of AI on the democratization of profit in financial markets?

To address this question, an agent-based simulation of the cryptocurrency market was developed, where we aimed to implement the scenario where a few key players of the market have access to an AI-powered tool able to predict the near future of the market and take conclusions in what comes to the impact on the returns of the common investor.

Literature Review

To study the scenario where the largest players of the market have access to advanced AI models that can accurately predict the future of the market, this work focused on building an agent-based model of the cryptocurrency market. In the context of financial markets, an agent-based model is a computer simulation where a given set of agents trade with each other on a strictly defined set of rules that aim to be representative of how markets work. The model usually follows an iterative structure, where a fixed set of steps is completed at each iteration. Only after having built a market simulation whose conclusions are scalable to real markets, can we implement the twist of the aforementioned hypothesis being a reality and compare the metrics of the market, in particular the impact on the democratization of profits.

The simulation being representative of how markets work is key. It is not doable to attempt to recreate every factor that goes into influencing the decision-making of investors. It would not only be impossible within a reasonable time frame, but most importantly, chances are that any observable result of such a complex simulation would be overfitted, that is, only applicable to the context of the simulation (Chen, Chang & Du, 2012). As such, when developing a simulation, two requirements stand out. The first one is to make it realistic, such that the conclusions of the object of the study apply to real life. The second one is to make it as simple as possible, so we can directly verify whether the mechanics and assumptions undertaken ensure this scalability.

This is the principle under which most authors since the beginning of the millennia have built agent-based simulations of financial markets, either as a basis to study a specific topic or with the only goal of creating a realistic market simulation. For this work, it is essential to understand how authors have come about materialising this principle within their studies, to guarantee the same is achieved in this one. Specifically, this boils down to understanding the rationale behind the assumptions taken on agent behaviour and market mechanics.

Starting with the latter, market mechanics refers to the way orders placed at each iteration by the agents - based on the pre-defined behaviour rules - are processed and matched, and how the market price reacts to them. The work by Raberto, Cincotti, Dose, Focardi & Marchesi (2005) defends the implementation of a limit order book. Following this mechanism, orders are stored in the order book until there is enough demand at a matching price on the opposite side of the market, at which point the orders are paired and a transaction is registered. This might happen immediately in the case of market orders, or after a few iterations for limit orders. The study concluded that this system results in far more realistic market metrics than using a cleaning house mechanism, where orders are filled immediately independently of supply or demand. Consequently, most authors in recent years have opted for this method for matching orders, and similarly, so will we with a few adjustments. In the work by Mizuta (2020), the author added a mechanism where orders expire and are cancelled after a given number of iterations. In what comes to the way the spot price is defined after each iteration of order matching, most studies take into account supply and demand to simulate how the price would naturally shift in a real-life scenario. In the simulation that Mizuta & Horie (2017) based their work on, the next spot price is set at the current intersection of supply and demand. In the work by Cont (2007), it is determined through a complex stochastic distribution that takes as a parameter the excess demand of the last trading round, a different alternative that follows the same principle.

Next, a review was conducted regarding how authors defined agent behaviour, which is the logic through which agents place orders. This includes the triggers that define when orders are placed, the rules that select the type of order placed (sell or buy, market or limit) and at which price and quantity. In this part, authors diverge more than they did in market mechanics. Nevertheless, comprehending the specific assumptions authors incorporated proves beneficial for determining the ones most pertinent to the considerations of this study.

Most authors create a small set of investor profiles, each having a distinct set of agent behaviour rules. In the work by Fratrič, Sileno, Klous & Engers (2022), agents were divided into randoms, speculative randoms and chartists. At each iteration, the first two types would hold the asset at scope with a given probability, which directly translates into buy or sell orders, with the only difference among them lying on the probabilistic distribution of the orders' limit prices. Chartist agents, on the other hand, would place buy or sell orders based on the trend of the asset's returns in recent iterations. In every case, the quantity of each order is defined by a random-distributed fraction of the agent's holdings. The study by Cocco, Tonelli & Marchesi (2019) also includes random traders who behave similarly; however, chartists decide when and how to trade based on the best trading rule chosen from a predefined set at each iteration by a complex genetic algorithm. In both cases, order prices derive from a fixed distribution centred around the current spot price. The work by Raberto, Cincotti, Focardi & Marchesi (2001) includes a more distinct set of traits. Firstly, every agent will trade at every iteration, and the position of buying or selling is picked randomly. While order prices also revolve around the current spot price, the actual stochastic distribution that defines them is set to make the expected value of the price for buy orders slightly higher than the spot price, and for sell orders slightly lower than the spot price. This is an option justified by the authors on the fact that traders placing orders would want to increase the chance of their orders being executed. Overall, this option validly increases market efficiency and is one we sought to implement in this study's model. Furthermore, the volatility of order prices increases proportionally to recent market volatility, aimed at considering the fact that when volatility is high, uncertainty about the appropriate price of an asset grows and so does the variability in the distribution of limit prices. Another work, targeted at studying hoard movements in financial markets (Cont, 2007), doesn't make a distinction about investor profiles. In this model, agents engage in trading based on their probability of participating in hoard movements.

This concludes the portion of the literature review dedicated to understanding the steps that are completed at each iteration of a broad range of studies. While authors diverge a lot in the specific rules that they put in place for agent behaviour and market mechanics, we can very clearly identify the prominent trait of attempting to recreate the basic components of markets, without going into much detail on every intricate detail that influences investor behaviour. Furthermore, whenever possible, authors try to follow rules that have been included in published works before, as generally these simulations serve only as a base for the true assessment of their work. Another homogeneous aspect is the structure of these simulations: iterative, where in each iteration agents start by placing orders based on their behaviour profile and then the simulation gathers these orders and makes the market function autonomously. This is the structure that was followed in this study's simulation, with rules on agent behaviour and market mechanics implemented with the same rationale as other authors have before, which will be further explored in the methodology section.

Before that, additional focus was directed towards other relevant parameters that authors set for their simulations which fall outside the scope of agent behaviour and market mechanics. For instance, a prevalent characteristic in these models is the inclusion of only one asset within the agents' scope, which serves as the sole option available for agents to trade. The only exceptions are rare cases where simulations include a range of assets that exist in real markets (Ankenbrand & Tomassiniy, 1997), which falls outside the scope of this study. Additionally, authors also abstain from features such as inflation or the continuous addition of more units of the financial asset throughout the market simulation, making the total amount of fiat currency and financial assets in the market time-invariant. Another relevant parameter is capital distribution, which translates into defining the amount of cash and asset holdings that each agent starts with. These metrics are updated throughout the simulation after each trade occurs and apply restrictions to how much agents can trade. In the work by Cocco L, Marchesi M (2016), the authors distribute crypto and fiat currency based on two distinct power-law distributions, bringing a realistic aspect into the simulation: that a small number of agents possess a significant portion of the total wealth. Others simplify this aspect by distributing cash and asset holdings uniformly across agents (Mizuta & Horie, 2017), while the author of the work Cont (2007) chose to stay out of this question by not keeping track of each agent's capital structure, removing the aforementioned restrictions.

With a comprehensive understanding of how authors have approached agent behaviour and market mechanics, as well as general market conditions, the next step is to transition into the methodology section, where it will be outlined the implementation of these components in this study's model.

Methodology

As shown in the literature review, successful agent-based financial market simulations tend to prioritize simplicity in their assumptions and market mechanics. The rationale behind this approach lies in the notion that while the simulation might not encompass every intricate detail of the market, the overall impact of any omitted specifics is, on average, reflected in the global market metrics and observations. Conversely, each simplifying assumption should accurately mirror the corresponding effect that all the intricate specifics of the market would exert. This is crucial in determining how these factors would affect the conclusions drawn had an ideal simulation of the market been developed.

With this in mind, the core principles of this model were built with a somewhat different approach to what other authors have done. The idea was to make the mechanic through which investors place buy and sell orders so natural that we don't need to take it as an assumption. But before that, let's address the overall structure of the model, and how we will proceed to detail every step taken to ensure both the simplicity and realism that were mentioned before.

As with every other study on this domain, the model takes on an iterative structure, meaning that the same set of steps are run at every iteration of the model. As such, to showcase the model we will go through every step that is completed in each iteration in normal market conditions and then through the extra components added to implement the scenario where some agents in the market have access to AI insights. But first, we will discuss the global market assumptions that were taken, followed by the procedure that is completed before this iterative process starts.

In this model, agents can trade one cryptocurrency for fiat currency. Throughout the simulation, no new units of cryptocurrency or cash are added to the model. This intentional exclusion of inflation and the absence of new agents joining the simulation ensure that the total amount of fiat currency and cryptocurrency tokens remains constant over time. Furthermore, the assumptions taken were made under the idea that each iteration of the model can be interpreted as a minute in real life, which will be reflected both in the length of the simulation and the time frame of the conclusions drawn.

Lastly, before the iterative process starts, agents are given a set of Fixed and Variable Attributes, which will influence their decision-making throughout the simulation. While Fixed Attributes are generated once before the simulation starts and are held constant, Variable Attributes also have their first instance generated before the simulation starts but, as their name suggests, might be changed under certain circumstances. Since the comprehension of these attributes depends on the context of the steps of the simulation, we will detail them in conjunction with their impact on each corresponding phase. The steps completed at each iteration are six-fold and are either steps aimed at recreating investor behaviour or tasks that ensure that the market mechanics run correctly.

The first of these steps determines which agents are set as active in the current iteration and was introduced to include in the simulation the realistic aspect that not every agent is active in every minute of the market. To account for differences in how active agents are, we assigned an 'Active Factor' as a Fixed Attribute, which determines how likely each agent is to be active on a given iteration. It was generated using a normal distribution clipped between valid probability values. The impact of an agent being active will be clear as the following building blocks of the model are explained.

Afterwards, for each active agent, the model moves on to update their 'Estimated Price'. This metric is the core of the model and one of the Variable Attributes. It is based on the idea that every investor, either consciously or not, has in their mind a price estimation of every asset that they are interested in. In real markets, this concept is influenced by an unmeasurable range of factors, including financial analytic calculations, but also news, sentiment found on social media posts, macro-economic factors, among many others. The rationale behind the inclusion of this factor in the model is that by building a reliable calculation of each agent's 'Estimated Price', we can directly derive agent behaviour in the form of the market positions that naturally arise from the comparison between this value and the spot price. The first instance of the 'Estimated Price' - calculated for every agent before the simulation begins - is a simple normal distribution centred on the arbitrarily determined initial spot price. Afterwards, at each iteration, active agents will update their 'Estimated Price' according to the following formula:

$$newEstimatedPrice \ = \ oldEstimatedPrice \ * (1 + k * \ \mathrm{N}(u, \sigma)$$

Where:

$$u = \ asset \text{ returns in the past } t \text{ minutes}$$

$$\sigma = \ asset \text{ volatility in the past } t \text{ minutes}$$

This formula includes not only the chartist behaviour as seen in the work by Fratrič, Sileno, Klous & Engers (2022), but also the mechanic seen in the study by Raberto, Cincotti, Focardi & Marchesi (2001), where the unpredictability of new estimated prices increases proportionally to recent market volatility. This formula also introduces two Fixed Attributes. The first one is the 'Market Influence Factor', represented as *k* in the formula. It is defined by a normal distribution and can be interpreted as how much an investor's perception of the asset's value is influenced by its recent performance. The second one is the 'Market Influence Time-Window', represented on the formula as *t*. Also defined by a normal distribution, it indicates the depth of the market analysis that each agent takes into account when predicting future asset prices. Both attributes were included in the model to create a continuous differentiation factor of how agents perceive the market to predict future returns. These not only increase the realism of the model but also avoid artificial patterns in the market that could appear as a result of a large set of agents behaving too similarly.

Then the model proceeds to define the positions that each active agent will want to take, based not only on the comparison between their 'Estimated Price' and the current spot price but also on their current 'Cash' and 'Crypto Holdings'. These attributes are two more Variable Attributes, once defined before the iteration process starts, and continuously updated based on the trades each agent makes. As in the work by Cocco L, Marchesi M (2016), 'Cash' and 'Crypto Holdings' are initially set out to resemble a power-law distribution, such that a small number of agents possess a significant portion of the total wealth, as it happens in real financial markets. As we will later see, these factors will limit each agent’s order magnitudes, and will be central to the end goal of the study. To bring this factor into the model, we distributed agent ‘Cash’ based on a Pareto distribution and derived the number of ‘Crypto Holdings’ that each agent starts with in a way that, based on the initial spot price, each agent’s capital in fiat currency equals the capital on crypto assets at the beginning of the simulation. This was implemented to negate any influence that initial capital distributions between cash and crypto could have and to avoid chronic supply and demand asymmetries which could result in prices rising or falling endlessly. Proceeding, in what comes to the position that agents are to take based on their ‘Estimated Price’ (EP), the core idea is that if an agent has an EP higher than the current spot price, they believe that the price of the asset should be higher than it currently is, and so they will naturally want to place a buy market order. Conversely, if an agent has an EP lower than the current spot price, they will want to place a sell market order, as they believe that the asset’s price should be lower. This logic is a robust way to create buy and sell orders through one single notion of EP. However, even though the differences in the quantity of buy and sell orders would alone shift the spot price, one essential mechanic is missing to ensure market efficiency: limit orders. Besides the orders placed at the current spot price, there should be times when agents would place limit orders at their current EP: buy limit orders if their EP is lower than the spot price and sell limit orders if their EP is higher than the spot price. To avoid having to implement forced criteria to define whether an agent is to place a market order, or the inverse limit order at their ‘Estimated Price’, we implemented the trait where both orders are made every time an agent is active. While this trait does not accurately resemble real investor behaviour, it is crucial to ensure realistic market efficiency, while avoiding undertaking more artificial rules based on arbitrary assumptions. To determine the quantities that make up each order, for market orders we implemented a mechanic such that order sizes are increased the larger the difference between an agent's EP and the current spot price is. This intends to render the magnitude of agent behaviour proportionally to the urgency that they would have to adjust their portfolio. Since this effect doesn't directly apply to limit orders, the quantities on these are a simple random proportion of an agent's current ‘Crypto Holdings’.

Next, it is time for the model to add these positions as orders in the order book. Like most authors in recent studies on this topic, we followed the limit order book structure suggested in Raberto, Cincotti, Dose, Focardi & Marchesi (2005). Essentially, two tables contain a list of all active orders, one for buy orders and another for sell orders. For each order, the tables keep track of the agent who placed it, the iteration *m* when it was placed and its price and quantity. The orders in the buy order book are arranged in descending order based on price, while those in the sell order book are organized using the same variable in ascending order. This arrangement ensures that the most competitive orders are positioned at the forefront. In instances where orders share the same price, they are further sorted based on *m* in ascending order, prioritizing orders placed earlier within the same price group.

Afterwards, the model moves on to the step where these orders are matched and fulfilled. To achieve this, the model fetches the first order from the table of buy orders (this process could be done entirely based on fetching the first order from the sell table and yield the same results) and checks if its price is higher or equal to the first sell order available. If so, we start iterating through the values on both tables, jumping to the next buy order every time one is fulfilled, and to the next sell order every time the current one can't entirely fill the buy order. At each match between a sell and a buy order, the transaction price is set at the middle point between the two prices, and as a new order is brought into view, we check if the buy order’s price is still higher or equal to the sell order’s price. The trading stops when this condition doesn't hold, or when one of the two tables is iterated through completely. This order matching system not only allows for multiple sell orders to complete a single buy order but also for orders to be partially filled if there are no more applicable opposite orders available. At this point, the model updates the ‘Cash’ and ‘Crypto Holdings’ attributes of each agent according to the orders that were partially or completely filled and removes from the order book the orders that were completely filled.

Lastly, we added a final step to maintain the order book. This step checks for any orders that were placed more than a given number of minutes ago and makes their prices slightly more competitive (increase the price for buy orders and reduce it for sell orders). If the new order price is higher than an agent's EP for buy orders, or lower for sell orders, the order is dropped from the order book, meaning that it has expired. This mechanism is crucial to make the market adopt realistic values of efficiency and represents the behaviour of agents allowing for less advantageous prices when the opposite demand/supply levels are not enough to match it, as long as it doesn't surpass their EP.

These steps complete a full iteration of the model running under regular market conditions, which brings us to the point of defining how the model simulates the influence of an Artificial Intelligence agent capable of predicting the near future of the market. But first, we need to define who should get access to it.

As mentioned in the introduction, we want to assess the scenario where this AI-generated intel is restricted to the most powerful players in the market. To simplify the application of this hypothesis, we limited access to these advanced insights to only one agent. This led to the creation of the last Fixed Attribute – ‘Has Access to AI’ – which in this case only has one positive value. The selected agent is always the one with the highest initial values of fiat currency (and consequently also of cryptocurrency holdings), which has its capital multiplied by a pre-defined factor to represent the financial leveraging power that this entity should have. Furthermore, the ‘Active Factor’ of this agent is set to 1, meaning that it will be active in every iteration of the market.

Now, in what comes to simulating the ability to predict asset prices in the near future, it is relevant to mention that since agent behaviour is stochastic, not even we know the future price points that the asset at scope will have. This means that the only way of simulating accurate future predictions is to reach the future and, based on the differences in market price through time, act as if the logical profitable positions were placed in the past. However, since our end goal is to estimate the impact of this agent on the profit opportunities for the common investor, it is fundamental that this agent's trading has a direct impact on the supply and demand mechanics of the market. This means that we cannot reach the future and act as if the correct position was taken with no regard to the supply and demand required to fulfil all the positions involved, which are themselves responsible for the determination of the market price. Taking all of this into account hereby follows the way we came about implementing this scenario into the simulation, which is directly aligned with the logic of an agent that is informed of the spot price of the market one minute into the future and acts accordingly: the AI agent is restricted from placing orders like the remainder of the agents do (based on their 'Estimated Price'). Instead, at every other iteration *m*, the model collects the positions from all other agents and obtains the market price that would come as a result of matching these orders. Then, we compare this "to-be market price" with the market price at iteration *m*-1, which can result in two inverse scenarios. To efficiently explain them both, we will detail the scenario where the price increases between iterations *m*-1 and *m*, followed between parenthesis by the scenario where the opposite takes place. If the 'to-be market price' is higher (lower) than the market price at iteration *m*-1, the model fetches every non-filled sell (buy) order that was in the order book at time *m*-1 with price lower (higher) than the "to-be market price" and performs the purchase (sale) of these orders at iteration *m*-1 by the AI agent. Then, we add to the current order book at iteration *m* a sell (buy) order at the "to-be market price" with the same quantity as the order at *m*-1 and re-run the order matching process. This method successfully simulates the scenario where the AI agent has access to the market price one minute into the future but relies on the remainder of the market to fulfil both the order in the present and the opposite one minute into the future.

This completes the methodology section, paving the way for the forthcoming section where the results derived from the simulation's execution will be presented.

Results

To obtain pertinent insights from the model, the simulation underwent 20 runs: 10 under normal market conditions - further referenced as Normal Market (NM) - and an additional 10 involving the utilization of the AI agent by a singular user with AI access – referenced as AI Market (AIM). Each of these 20 runs encompassed 10 080 iterations. Given that each iteration represents a minute in reality, each simulation run simulates a week in real-world time. According to the characteristics of the model, this length was sufficient to include all the relevant insights that we could extract from it. In every run there were 10 000 agents trading in the market (one of them having access to AI-generated intel in AIMs). This is a population size regularly used by authors in agent-based simulations of financial markets, and one that in this case resulted in satisfactory market dynamics.

Let's start by examining the difference in the patterns that emerge from the effect of agent actions in NMs and AIMs in the evolution of market prices. In the left-hand figure below, we observe an example chart depicting market price fluctuations over a 6-hour period under NM conditions, representative of how price usually varies under this scenario. Here, the price undergoes smooth changes, fluctuating a few percentage points both upwards and downwards, devoid of significant spikes. Contrastingly, the right-hand chart is representative of the impact of having an entity benefiting from AI predictions on market prices. In this scenario, prices exhibit more pronounced and erratic shifts due to the AI-driven agent influencing demand and strategically leveraging positions to yield profits in the forthcoming iteration.

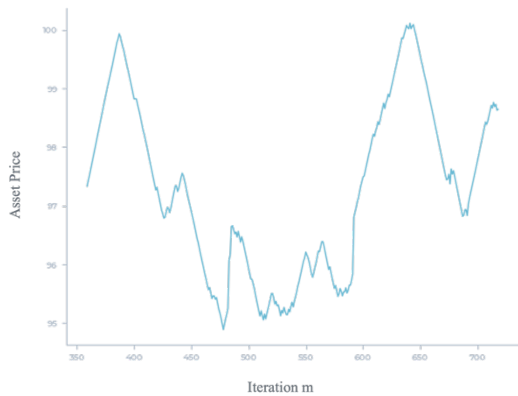

Figure 1A: Example representative of asset price fluctuations under NM conditions

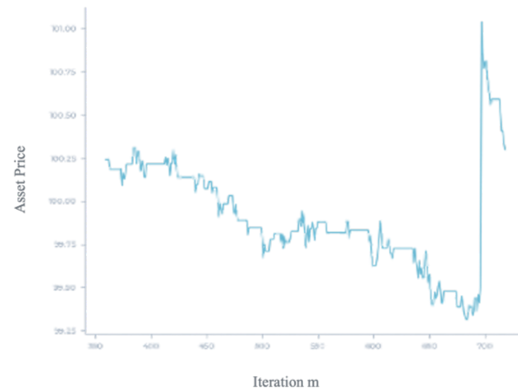

Figure 1B: Example representative of asset price fluctuations under AIM conditions

Nevertheless, when observed over a weekly scale, price fluctuation charts look relatively similar between NMs and AIMs, as the distinct price disparities become less noticeable. Despite the differences in short-term price movements, both market conditions tend to maintain prices within a narrow range of around a few percentage points of the initial spot price. This stability arises from the absence of significant hoard movements typically observed in real markets.

Moving forward, the evolution of the distribution of estimated prices is also similar in the two scenarios, presenting comparable values of the standard deviation of estimated prices across time. Both scenarios demonstrate a slight increase in these deviations as the simulations progress. A relevant metric for evaluating the reliance on the notion of estimated price can be obtained through the percentage of agents that, at different stages in the simulation, experienced their estimated price exceeding the spot price at one point and falling below it at another point. This metric exceeds 99.9% in both market scenarios.

Disparities start to arise as we move on to the most important aspect of the results: how did the introduction of the AI agent in the market affect the returns of regular investors? When adjusted for the initial spot price, normal investors averaged 0.96% returns under the span of one week on NM conditions. This value decreased to -0.39% after the introduction of the AI agent. This means that under this AIM, regular investors are expected to lose capital by entering the market while on NM conditions the case was the opposite. The AI agent, on the other hand, was able to average 9.64% returns within a week. This is further demonstrated in the two histograms below, depicting the distributions of normal investor returns under the NM (on the left) and the AIM scenario (on the right). While both distributions present similar values of kurtosis, the latter is slightly skewed towards the left, which is representative of this averaged decline.

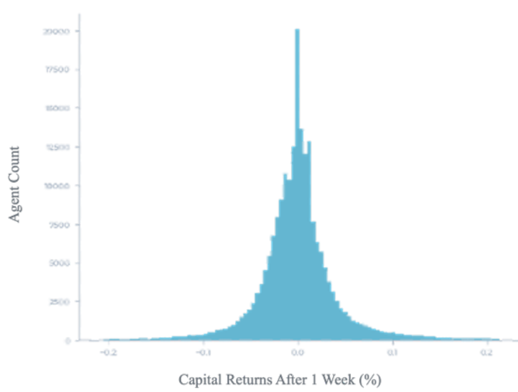

Figure 2A: Distribution of investor capital returns over one week for 10 NM iterations

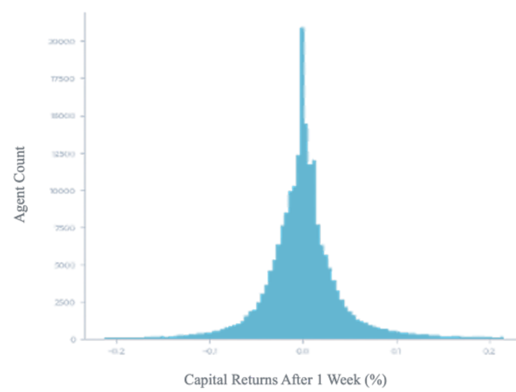

Figure 2B: Distribution of investor capital returns over one week for 10 AIM iterations

Another symptom of the impact of the AI agent can be found in the size of the buy and sell order books. On NM conditions, both order books average around 700 orders on each iteration, with a standard deviation of around 110. This means that close to 7% of the agents have, at each point in time, an active order waiting to be fulfilled. This results in healthy levels of efficiency and volume throughout the simulations, further suggesting the successful implementation of the mechanic based on 'Estimated Price'. On AIMs, however, this size averages around 500 in the first iterations and quickly drops to an average of 100 orders after an artificial length of an hour. This level of order scarcity remains fairly consistent throughout the remainder of the simulation and is the result of the AI agent leveraging the combination of its financial power and the AI insights that it receives to fill every “wrong” order on the market to profit in the upcoming iteration. This effect can also be felt on the percentage of orders that expire. While on NMs around 14% of orders expire, in AIMs this value drops close to 6%. Since these orders are mostly being filled to fuel the profits of the AI agent, we could say that generally it would have been better for normal investors if these orders had remained untouched.

Conclusion

Throughout this study, the simulation presented compelling evidence regarding the potential ramifications of introducing AI agents capable of predicting market prices into financial ecosystems. The contrasting regular investor returns before and after the AI agent was introduced underscore a significant shift in market dynamics. Regular investors’ expected returns turned negative, while the AI agent consistently achieved notably high returns within the simulated environment.

These findings raise pertinent questions about the future landscape of financial markets. As AI-equipped entities gain an edge, there emerges a concern for the equitable participation of regular investors. The scenario simulated in this study, while a hypothetical construct, foreshadows potential challenges for those without access to similar AI tools, possibly dampening confidence or involvement in financial markets among individual investors.

Looking ahead, it is imperative to engage in discussions and explore measures that promote fairness, transparency, and accessibility within financial markets amidst the increasing adoption of advanced AI tools by influential market entities. Ethical considerations, regulatory frameworks, and the democratization of AI access in finance demand further attention and deliberation.

While this simulation offers valuable insights, it's essential to acknowledge its limitations and areas for future research. Further refinements in modelling and exploration of nuanced scenarios can provide a more comprehensive understanding of how AI influences financial markets and its subsequent impact on individual investors.

References

1. Raberto, M., Cincotti, S., Dose, C., Focardi, S. M., & Marchesi, M. (2005). "Price Formation in an Artificial Market: Limit Order Book Versus Matching of Supply and Demand." In T. Lux, E. Samanidou, & S. Reitz (Eds.), Nonlinear Dynamics and Heterogeneous Interacting Agents (Lecture Notes in Economics and Mathematical Systems, Vol. 550). Berlin, Heidelberg: Springer. Retrieved January 1, 2005, from https://doi.org/10.1007/3-540-27296-8_20
2. Mizuta, T., & Horie, S. (2017). "Why do active funds that trade infrequently make a market more efficient? — Investigation using agent-based model." 2017 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1-8. Honolulu, HI, USA. doi: 10.1109/SSCI.2017.8280798
3. Cont, R. (2007). "Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models." In G. Teyssière & A. P. Kirman (Eds.), Long Memory in Economics. Berlin, Heidelberg: Springer. Retrieved January 1, 2007, from https://doi.org/10.1007/978-3-540-34625-8_10
4. Fratrič, P., Sileno, G., Klous, S., & Engers, T. (2022). "Manipulation of the Bitcoin Market: An Agent-Based Study." Financial Innovation, 8(1), 1-29. Springer; Southwestern University of Finance and Economics.
5. Cocco, L., Tonelli, R., & Marchesi, M. (2019). "An Agent Based Model to Analyze the Bitcoin Mining Activity and a Comparison with the Gold Mining Industry." Future Internet, 11(1), 8. Retrieved January 1, 2019, from https://doi.org/10.3390/fi11010008
6. Raberto, M., Cincotti, S., Focardi, S. M., et al. (2003). "Traders' Long-Run Wealth in an Artificial Financial Market." Computational Economics, 22, 255–272. https://doi.org/10.1023/A:1026146100090

7. Chen, S., Chang, C., & Du, Y. (2012). "Agent-based economic models and econometrics." The Knowledge Engineering Review, 27(2), 187-219. doi:10.1017/S0269888912000136

8. Mizuta, T. (2020). "An agent-based model for designing a financial market that works well." 2020 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 400-406. Canberra, ACT, Australia. doi: 10.1109/SSCI47803.2020.9308376

9. Cocco, L., & Marchesi, M. (2016). "Modeling and Simulation of the Economics of Mining in the Bitcoin Market." PloS One, 11(10), e0164603. https://doi.org/10.1371/journal.pone.0164603

10. Ankenbrand, T., & Tomassini, M. (1997). "Agent-Based Simulation of Multiple Financial Markets." Presented at PASE'97, November 9-12, 1997.