International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 2, Issue 3, 21-29, 2021
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V2I3P103

Original Article

A Multi-Agent Approach to Market Microstructure Modeling: AI Perspectives on Financial Liquidity

Archana Pattabhi
Executive Leader in AI, Cybersecurity & Risk, SVP Citi; Member, Forbes Technology Council; CIO/CISO Advisory Board, The Executive Initiative.

Abstract *- Some considerations of the working of the financial markets form part of the basis of theoretical and practical finance. Market microstructure on how securities are traded, and prices are determined has been impacted on by algorithmic trading and Artificial Intelligence (AI). In this paper, a new approach to modeling the financial liquidity of a company is suggested, characterized by using a Multi-Agent System prototype with artificial intelligence built inside the acting agents. It will assess how heterogeneous agents, through their behavior and interaction, impact the aspects of market liquidity, bid-ask spreads, volatilities, and price discovery. The current paper proposes the theory of using intelligent agents, RL, and supervised learning to incorporate institutional traders, market makers, and retail investors. The proposed work involves implementing a simulation environment that models a LOB and uses real market data for evaluation. Pursuant to this, the simulations reveal that the AI-enhanced multi-agent models provide a near-perfect simulation of the actual markets and, more importantly, provide a better understanding of the microstructure of the markets in relation to information asymmetry and order strategies of the various players in the market. Therefore, agent-based decisions employing neural network-based decision functions show complexity, such as clustering, flash crashes, and temporary illiquidity. They are examined via several market indicators and representations. Furthermore, it is evaluated under various regulatory and market conditions, such as circuit breakers and liquidity injections, to derive policies. This paper's scholarly contribution would be an attempt to move the debate beyond the current vehicle of analysis that is predominantly constituted by economic theories and frameworks and seek the help of AI techniques in doing so.*

Keywords *- Market Microstructure, Artificial Intelligence, Multi-Agent Systems, Financial Liquidity, Reinforcement Learning, Limit Order Book, Algorithmic Trading, Price Discovery, Simulation, Deep Learning.*

1. Introduction

1.1 Background

On this background, the microstructure of any financial market is of profound importance in making economic decisions. It would capture trading processes, actors, actions, and the execution of orders on the trading platform. [1-4] In the last few years, the improvements in high-frequency and algorithmic trading systems have completely altered the conventional microstructure models.

1.2 Role of AI and Multi-Agent Systems

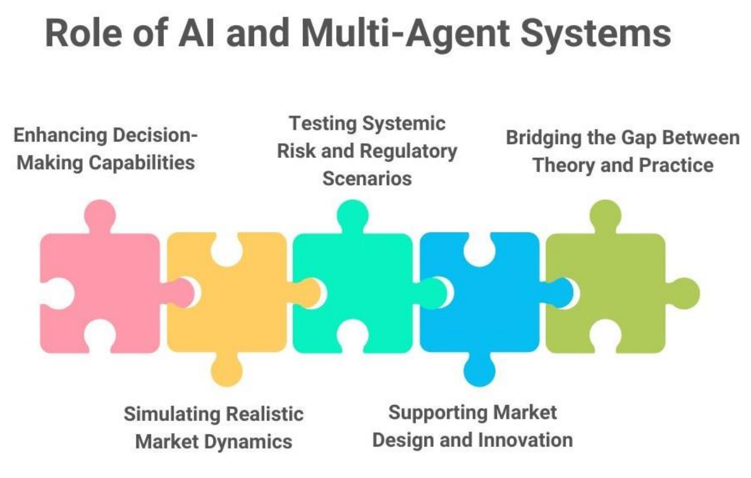

Fig 1:Role of AI and Multi-Agent Systems

- **Enhancing Decision-Making Capabilities**
  In financial markets, AI has changed the execution of decisions by ensuring that agents draw experience from realistic data, apply it to changing circumstances, and then align it to seek better solutions. In contrast to the rule-governed rationality approach to decision-making in agent systems, machine learning, and reinforcement learning facilitate the adaptive behavior of the agents according to the signals coming from the market environment. This capability is critically used in real-time, especially in high-frequency and algorithmic trading areas where timing and accuracy are determinative.
- **Simulating Realistic Market Dynamics**
  Multi-agent systems (MAS) provide a paradigm for describing a diverse global environment where stakeholders have various objectives, available means, and behaviors. MAS incorporates all the characteristics of a live market through the replication of market makers, institutional investors, and retail traders. They become endogenous depending on the time horizon, and it is possible to observe already-known economic phenomena: liquidity cycles, volatility clustering, and price formation. This makes it possible for researchers to analyze the resulting emergent behavior from local decisions with interaction from agents within the system.
- **Testing Systemic Risk and Regulatory Scenarios**
  Thus, one of the most important applications of AI in multi-agent simulations is its use to model the market's reactions to the maximum and stress situations. To investigate these aspects, scholars have to create external shocks, such as sudden liquidity pull-out or regulation changes, to help them see how these affect the AI agents and how their reactions ripple through the market. This has great benefits for regulators because it enables them to assess the risks that some patterns of conduct pose for the whole financial system and experiment with measures that will be adopted in the actual markets.
- **Supporting Market Design and Innovation**
  Together, AI and MAS act as an environment where novelties in terms of market designs and structures, tradingmethodologies, and auctions can be introduced and analyzed. Analysts can model what would happen if one introduced new instruments, a new fee structure, or changed the rules in executing certain trades. As such, they simulate how real market actors would behave, thus making AI agents valuable sources of inspiration for the growth of smarter markets, especially for exchanges, banks, and even policymakers.
- **Bridging the Gap Between Theory and Practice**
  Certain assumptions are taken into economic models that may not be very practical, thus causing the emergence of new economic models. Multi-agent systems integrated with AI afford the connection of these frameworks and the actual strategies seen in today's markets. The reinforcement of economic principles within adaptive, learning-based agents gives analytical mechanisms in such systems a more realistic and empirically grounded qualitative. This approach allows for posing questions that are computationally difficult and also for testing models against data.

1.3 Emergence of AI in Financial Markets

In recent years, integrating artificial intelligence (AI) into the financial markets has brought radical changes in trading, risk management, and analysis of the financial markets. Historically, there were conventional theories that originated from economic principles, which, however, used linear thinking and strategies that could not well explain the complexities involved in the current financial dynamics. In this context, they see that Machine learning (ML) and Reinforcement Learning (RL) bring a new perspective that can be described as more intelligent modeling to the financial field. [5,6] These methods enable algorithms to derive, from large and real-time data, historical and new trends and then decide, to some extent, independently about future actions. These include tasks such as price forecasting, sentiment analysis, fraud detection, and portfolio optimization, to mention but a few. Different methods like Support Vector Machines, Random Forests, or deep learning techniques like LSTM and CNNs can be used in the high dimensional financial data to extract useful features of datasets and the dependencies between them.

It has enhanced the accuracy and timeliness of business forecasts and territorial differentiation in trading systems. The former is extended in reinforcement learning, where agents can develop the best strategies to tackle the same environments through practice. Some algorithms that can be employed for designing autonomous trading agents include Q learning, Deep Q learning, PPO, and other strategies where the trading agent can learn from the market feedback and modify the behaviour to achieve long-term rewards. They are realistic models that mimic decision-making in uncertainty; they combine risk-taking by identifying new opportunities with taking advantage of rewarding actions. A combination of ML and RL has become a way of creating smart agents with the sensibility to respond to the market and the ability to influence it. As these technology advances, financial markets become adaptive ecosystems where data intelligence is one of the key factors for success.

2. Literature Survey

2.1 Traditional Market Microstructure Models

The market microstructure theories have long dominated the analysis of trading and the mechanics of price determination in securities markets. The Glosten-Milgrom approach is a model that presents the issue of market makers who post prices when there is information asymmetry in the form of informed traders. Conversely, Kyle (1985) explores a broad, rational, informed trader who disseminates the information he holds through trading to make profits without distorting the asset's price. [7-10] These models emphasize the importance of information asymmetry, order flow, and liquidity, which present the basic framework through which information is processed and how prices are formed.

2.2 Agent-Based Modeling in Economics

Agent-Based Modeling (ABM) has become widespread in economic and financial simulations since it makes it possible to introduce heterogeneity among the agents and their learning algorithms. Lux and Marchesi (1999) presented one of the first ABMs in finance to show how the interaction of fundamentalist and charist agents can produce such phenomena as volatility clustering and the power-law distribution of returns. Due to their handy method for expressing the agents' standing micro behavior and micro-interaction, ABMs can be especially helpful in demonstrating the emergent properties of the global markets. Thus, they can be useful in analyzing the complex systems in the economy.

2.3 Machine Learning in Market Analysis

ML has benefited a lot in market analysis since it brings a different approach to analyzing high-dimensional data and making conclusions. Current methods include support vector machines, decision trees, and some of the state-of-the-art approaches, such as LSTM and CNN. These are trends analysis, risk assessment and analysis, and algorithm trades. In such an ever-changing and complex environment, where many factors may interact in unknown ways, the strength of ML models to identify nonlinear relationships and learn from large historical data makes them useful tools for analysis.

2.4 Reinforcement Learning in Trading Agents

Reinforcement Learning (RL) has recently become more popular due to its proposed application in dynamic environments such as trading. While in the case of supervised learning, the comprehensible data previously tagged with desired outputs are given to the model, in the case of RL, the agents learn from interactions with the environment and get feedback regarding rewards. These include Q-learning, DDPG, and PPO, a few strategies for designing adaptive and robust trading agents for financial market simulation. These features allow the agents to adapt their behavior according to market changes and find optimal and often unconscious trading patterns, which can produce promising directions for automated trading systems.

3. Methodology

3.1 Agent Design

The elements of the proposed agent-based market model [11-14] are agents that represent various kinds of participants in a market, with actions and goals.

Fig 2: Agent Design

- **Market Makers (MM)**
  Market makers must supply continuous quotes of bid and/or ask prices to ensure an active market. Liquidity providers' main goal is to earn a profit based on the bid-ask spread price difference and control the risks arising from the stocks they hold in their possession. They are always quick to respond to the flow of orders and the market

situation to quickly stabilize the price differences. In general, simulations of market makers apply algorithms using quotations depending on the supply-demand and price movement.

- **Institutional Traders (Inst)**
  Institutional traders are the major traders, including hedge funds, pension funds, mutual funds, and any other substantial traders in the foreign exchange markets. The first and foremost purpose is to avoid affecting the particular market and trade expense as much as possible while attaining the best results. These agents use what we know as execution algorithms such as VWAP, which stands for Volume Weighted Average Price, or TWAP, which stands for Time Weighted Average Price; they ensure that they split the large order into a number of working orders spread over time to minimize on the slippage risks, or to avoid causing a screw on the market.
- **Retail Investors (Ret)**
  Keith is a small individual investor, also known as a retail investor, who usually trades with relatively basic trading strategies. They could be momentum trading strategies that buy when prices are high and sell when prices are low or mean reversion strategies that believe prices will return to their mean. Thus, retail agents tend to have limited resources available to them and are inherently myopic traders. Their trades may even contribute to the noise in prices and the trade volume.

3.2 AI Decision Engines

As for credibility and adaptability to the market and other agents, each type of agent has an artificial intelligence AI decision-making system based on its goals and scale of operation. These engines are based on the built-up of neural networks to implement agents' learning nature and growth. Market makers involved in maintaining market depth and inventory risk are Market Makers use a particular type of learning called reinforcement learning that aims to make the best choices based on value function. The market environment includes information on order flow and price changes within the recent period, and the Q-learning agent deploys actions that it expects to reap high future utility, such as crossing the bid-ask spreads while incurring the low penalty of holding an unbalanced inventory. The Q-function is estimated by a neural network, which allows the agent to make reasonable predictions for the other types of similar market states. By the same token, retail investors are modeled using supervised learning classifiers since they rely primarily on rule-based approaches and are less complex than arbitraged decision-making.

These agents learn from past observations in the market to categorise trading signals such as buy, sell, or hold using some of the indicators such as moving average, momentum, or price patterns. This is similar to general naive trading behavior when they chart via line analysis techniques or other simple heuristics. Altogether, these AI-based modules offer a wide and realistic range of the agent’s actions in the simulated market. Using reinforcement learning, institutional traders who transact in large quantities of contracts and require advanced trading techniques constitute a trading agent known as the Deep Deterministic Policy Gradient (DDPG). Due to the continuous action space, DDDPG is well-suited for fine-grained order placement and timing decisions. The institutional agent’s DDPG has actor-critic networks that allow the agent to learn a deterministic strategy and a value function that can maximize trades in terms of time while incorporating the effect of the impact cost and transaction cost. Such a setup enables institutional agents to work under dynamic conditions to be efficient in execution.

3.3 Simulation Environment

The most important structural elements of the simulation are contained within a Limit Order Book (LOB), which is a record of buy and sell orders placed by artificial agents. It mimics a market akin to the trading exchange since it accepts orders, saves them, and executes them in specific ways. It is in this case that agents can place two major order types; these include limit orders, which indicate the price below or above which the agent is willing to purchase or sell an asset, and the market order, which is an order that is placed and executed at the best available price. [15-18] The LOB also enforces the price-time priority where higher bid prices for the buyers are given preference and the lower ask prices for the sellers, and among all the orders with equal prices, the order that has been entered first is programmed to be executed. This makes the environment very dynamic. Kant's menu is expensive in terms of the nature, size, direction, type, and time of placing an order, all of which are strategic capabilities the agents need to consider.

LOB reflects the current state of orders and instructions and gives information on the depth of the market by showing the bid and asking price queues. Market makers employ the information in controlling the spread and the state of their book, while institutional traders evaluate the degree of liquidity and price blunting or slippage for themselves and their clients: retail agents act on market indicators or signals. This time is discrete, and during this period, agents see the current state of the order book and then submit their orders by running their AI strategy. Prices for the transactions are arrived at from matched orders, and all unmatched orders are held in the book until canceled or executed. This specific and highly regulated LOB simulation engulfs key market elements like supply and demand, price formation, and volatility. Thus, it can be best used to test how different financial representatives work worldwide.

3.4 Market Scenarios

In order to assess the stability and flexibility of the agent behaviors, several different market settings are tested, representing rather realistic market situations and stress factors.

Fig 3: Market Scenarios

- **Normal Market**
  In normal market conditions, trading circumstances are regular, so the orders are continuously appearing and being processed, supply and demand for the instruments are moderate, and volatility is within regular range. This environment is still simple and can be used to set a baseline for managing the agents' performance under standard circumstances. It depicts a market where trading is active, costs are narrow, and prices are determined seamlessly.
- **High Volatility**
  Specifically, the high volatility causes wide price changes and higher risk levels, which are usually caused by the release of information events or changes in the macroeconomic environment. In such circumstances, the order flow gets disrupted, the spread decomposes, and the possibility of slippage rises. This challenges the agents to effectively negotiate risk factors, think fast, and work efficiently even in an uncertain environment. It also puts reinforcement learning agents in a position to update their value functions more often because of quickly changing market states.
- **Sudden Liquidity Withdrawal**
  This influence is modeled by an abrupt stop-out regime where most liquidity providers (such as market makers) pull out of the market en mass. Again, this reduces the liquidity available, moderated by high bid-ask spreads, low volumes of trades, and possibly generated price gaps. It emphasizes the robustness of the trading strategies and the ability of the remaining brokers to accommodate or offer the stabilizing forces in an illiquidity state.
- **Regulatory Intervention**
  In this case, the rules used externally apply mid-simulation; these include the transaction taxes, the order throttling, and the price limits, also known as the circuit breakers. The idea is to see how agents react to regulation patterns and whether it is indeed possible to bring more stability to the markets or, on the contrary, such measures harm the trading environment. This paper provides an understanding of politeness from an applied perspective. It analyzes how the identified policies enable us to understand the effectiveness of certain policies from a strategic perspective based on the changes in rules.

3.5 Metrics for Evaluation

The variables in this selection represent a fairly objective set of features for evaluating the microstructure from the point of view of quantitative characteristics. [19,20] It established metrics encompassing some of the fundamental parameters of market quality, efficiency, and stability.

- **Liquidity Depth**
  Liquidity depth refers to the ability to handle the number of limit orders within a range a few ticks away from the mid-price, the current bid-ask midpoint. This means that the larger the size of a book, the more volatility the market lacks when big trades are made. It is an excellent representation of the level of openness of the positions for participants. It is more significant for institutions like traders and market makers, given their focus on execution risk.
- **Bid-Ask Spread**
  This is a measure of the difference between the highest price a buyer could hope to secure stock for (known as the bid) and the lowest price a seller could hope to get for his shares (referred to as the ask). A small spread means many competitors in the market, and transaction costs are low. Conversely, a high spread means that there are few players in the market or some uncertainty about the price that reflects the asset's real value. This point is useful for assessing the effect of the agent’s activities and the presence of various factors on trading costs and the efficiency of prices.

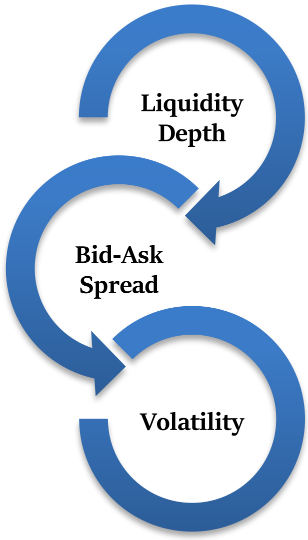

Fig 4:Metrics for Evaluation

- **Volatility**
  Volatility is defined as the sum of the standard deviation of price returns for a certain time period and, in general, is used to measure the variability of the prices of a certain instrument. Higher volatility offers the traders a higher risk and thus could lead to very volatile situations in the trader's trading arena; on the other hand, lower volatility implies that the market is stable. This measure is important in determining the risk map of some situations and actions of agents, and it is sensitive to stressful situations such as increasing volatility or outright withdrawal of liquidity.

4. Results and Discussion

4.1 Liquidity Under Normal Conditions

The state of the markets in normal conditions was proven to be rather stable and efficient, as it conforms to various characteristics of financial markets existing in the real world. Market makers deserve special praise for keeping the market relatively liquid by submitting new buy and sell limit orders on the appropriate side of the middle price. They moved their quotes flexibly with changes in the flow of orders and movements of the market, keeping a narrow bid-offer spread that helped to minimize costs for all the agents involved. They constantly maintained their limit order book, which was the most helpful in ensuring proper continuity and steady liquidity provision to other agents on the platform. Larger order flows of institutional traders and their trades' sophisticated execution strategies explained the depth of the market. This they did, spread out over time and price levels, helped to reduce market impact and improve the order book in general. Their actions may also maintain the stability of the price fluctuations for the short-term volatilities of big buy or sell pressure without necessarily creating volatility.

This interaction between the institutional and market-making agents provided a strong two-sided market with proper provision and consumption of liquidity. Thus, retail investors, less strategic but fewer in numbers and volumes, supplemented trading behavior diversification. Most of them used far less sophisticated strategies, like momentum or mean-reversion, thus adding a certain level of variability to mimic the real market environment. Their trades ensured that the market was active and, most essentially, magnified the trends regarding the outlook of their sentiment, which made the simulation highly behavioral. Altogether, the well-coordinated working of all the agent types under normal circumstances indicated the high degree of order matching, the stability in price fluctuations, and the low extent of slippage in the trading. The K-line had a reasonably dark book at the mid-price, and there were no significant order imbalances during the formation of the price resetting process, also known as price formation.

4.2 Impact of AI Decision-Making

Applying the adaptation engines based on the RL, main advances in the agent's behavior and the general market performance were observed. It was distinguishing that AI-based agents were not program-based like rule-based agents. They devised different strategies and modified them based on their experiences in the simulated market. This was most vivid when testing was conducted on learning models such as DDPG and Q-learning models, outperforming even the standard ones. These agents were able to learn specific and multiple aspects of trading and control, including the size of the order, the timing of the placement of the order, and the management of the spread in relation to certain signals from the market, such as volatility, flow of orders and depth of the market. This improvement is depicted in Figure 3, where new liquidity measures of average bid-ask spread and order book depth are presented for both AI-based and rule-based agents. Market parameters such as average spread and depth were also improved, narrowing down to the top five price levels confirmed by a decline in the transaction cost of the AI-enhanced market.

This performance increase was due to the fact the AI agents could ''see'' the dynamics of the order flow and the near-term movements of the markets and be able to act on this information faster than any conventional rule-based systems. An account of spearheading a highly crucial sector of the strategies involved in market-making, inventory was also well-handled by AI agents. Ensuring that the profit opportunities are always balanced against the risks arising from inventory fluctuations, these agents could maintain a steady place in the order book and move dutifully whenever there are minor fluctuations in volume. This approach also reduced the issue of adverse selection. It ensured that AI agents could take advantage of short-term inefficiencies in the markets while helping maintain market liquidity. Therefore, the enhancement of agents' capabilities through reinforcement learning allowed them to make better decisions informed by the current market condition, which contributed to effectively maintaining the market stability, decreasing costs for participants, and being resistant to fluctuations in the market.

4.3 Emergent Behaviors

One of the most significant observations pertaining to the simulation was that new and sophisticated market patterns arose that were not incorporated as plans of action by the agents. These complex emergent phenomena arise from simple rules or AI-based heuristics embedded in a common and simulated marketplace and mimic actual markets.

Table 1: Percentage Improvement of AI-Based Agents Over Rule-Based Agents

Metric	Improvement
Avg. Spread	41.67%
Depth	40.00%
Volatility	21.74%

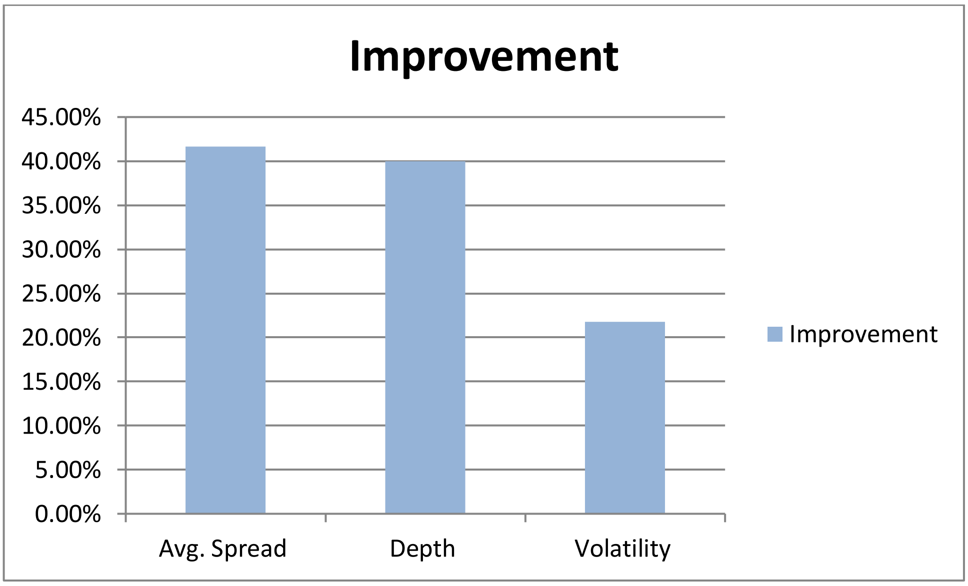

Fig 5: Graph representing Percentage Improvement of AI-Based Agents Over Rule-Based Agents

- **Flash Crash Events**
  A flash crash-like situation was observed when many agents resorted to a similar behavior to how they responded to cues or volatility in a continuously changing market environment and started selling in unison. This selling, coupled with market makers pulling out their funds from the market quickly, also triggered a sharp decline in its prices. Even though the market returned to normal after some time, it demonstrated the main points of the real-life flash crashes and the weakness of the market conditions for collective actions and low market liquidity.
- **Trade Clustering**
  Another observed emergent behavior included the tendencies of the trades to group themselves in relation to psychological or technical prices. Agents most likely to submit orders were the retail and momentum ones, specifically entering orders close to round numbers and previous highs and lows. This gave further momentum to these levels and strengthened short-term activities while generating occasional temporary disturbances to the thin volume levels.
- **Feedback Loops and Volatility Bursts**
  Feedback loops were also used, particularly by the agents who employed momentum-based incentives. This would increase the price as more trend-following agents placed more buy orders, consequently increasing the price. These virtuous cycles, in turn, caused short-lived volatility, even if no external factors led to such fluctuation. That being the

case, such dynamics shed some light on the realism and usefulness of agent-based modeling in capturing systemic risk and market vulnerability.

4.4 Market Stress Scenarios

In situations of high-risk conditions, for instance, withdrawal of liquidity, volatility, or sudden appearance of regulations, the AI-based agents outperformed the rule-based agents by a big margin. Stress testing was all the more important in the agent-based market as it allowed us to determine how any given trading strategy would work in an unstable, uncertain, and dynamically changing environment that was typical for real-world markets. The rule-based agents, specifically, were slow or inapplicable to the changing market states due to their inherent structure, which was rigid, logical, and contained threshold values. For instance, during a condition of liquidity shock, the market makers disappeared from the list, and rule-based agents persisted with placing orders, ignoring any extra risk it incurred a negative effect triggering inefficient trades, giant variation in bid and offer, and disseminating effect on price changes. On the other hand, the AI-based agents, specifically those shed with reinforcement learning, such as DDPG and Q-learning, did a much better job detecting changes in the market's movement and responding to them.

They adaptively adjusted their order arrival timing, bid-ask prices, and risk-taking behaviour vis-à-vis stress indicators, including variability, depth or shifts in the order flow. For example, in volatile time periods, the AI agents increased their quoted spreads to cover risks while remaining active in the market to make it functional and liquid. They also made precautionary movements, reducing order size or moving to a safer price level before an event. This adaptability will be investigated by comparing order flow and the spread between AI and rule-based agents in a simulated liquidity crisis. Whereas the rule-based agents aggravated the price shocks with lesser and delayed response time, the AI-based agents acted as a buffer, thereby softening the intensity and the disruption time span. These findings shed light on how AI can enhance market robustness and can be important for institutions in charge of regulating the markets and for the exchanges interested in implementing intelligent trading solutions in real-life financial markets.

5. Conclusion

This work demonstrates that incorporating AI in multi-agent methodologies provides the basis for a versatile machine that can effectively emulate the complexity of the financial markets. With the help of the agents for the implementation of this research, such as market makers, institutional traders, and a crowd of other kinds of agents, which are the retail clients, the simulation framework could capture numerous emergent behaviours, the fluctuations of liquidity as well as the stress responses of the financial market that has some parallels in the real market. There was also an integration of artificial intelligence, especially reinforcement learning, to allow agents to make contextual decisions and adjust to their environment and market changes quickly. This level of adaptability bore fruits in parameters, that is, bid-ask spreads, depth, and volatility levels, especially during thin trading periods or high volatility as a result of withdrawal of liquidity. The sophisticated patterns like flash crashes or clustering of trades and feedback-driven volatility provided evidence for the richness of the model. These behaviors were not predefined but emerged from the agent's characteristics concerning their risk-taking propensity, trading strategies, and learning abilities. It becomes useful for financial regulators and other institutional players since it presents such scenarios as a place to experiment with real-world market settings and evaluate the impact and efficiency of interventions or alterations in the system.

From a pragmatic perspective, the result of this research is beneficial as it shows how AI can be used to improve the flow of markets. As AI agents demonstrated the ability to rebalance when trading volumes are high, spreads increase, or adjust their position in the inventory, it also proves that intelligent trading can help stabilize the market during volatile situations. This raises questions about the use of algorithmic trading in reducing systematic risk and the future of AI regulation in the financial sector. As for further development, it is possible to point out quite a number of directions in the further development of the given framework. Uniquely, such aspects as multi-market interactions, cross-asset correlation dynamics, and different inter-agent communication topologies and protocols can be added to further expand the model’s realism. Finally, adopting more complex architectures like transformers, GNN, or both can also help enhance the ability of agents to make decisions. With the increase of complexity in the financial markets comes a wealth of opportunity for the endless advancement of how AI, economics, and compounding computational modeling can be used to create new, stronger, efficient, and intuitive financial systems.

References

[1] Van der Hoek, W., & Wooldridge, M. (2008). Multi-agent systems. Foundations of Artificial Intelligence, 3, 887-928.

[2] Ferber, J., & Weiss, G. (1999). Multi-agent systems: an introduction to distributed artificial intelligence (Vol. 1). Reading: Addison-wesley.

[3] Uhrmacher, A. M., & Weyns, D. (Eds.). (2018). Multi-Agent systems: Simulation and applications. CRC press.

[4] Grzonka, D., Jakóbik, A., Kołodziej, J., & Pllana, S. (2018). Using a multi-agent system and artificial intelligence to monitor and improve cloud performance and security. Future generation computer systems, 86, 1106-1117.

[5] Tweedale, J., Ichalkaranje, N., Sioutis, C., Jarvis, B., Consoli, A., & Phillips-Wren, G. (2007). Innovations in multi-agent systems. Journal of Network and Computer Applications, 30(3), 1089-1115.

[6] Glosten, L. R., & Milgrom, P. R. (1985). Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. Journal of financial economics, 14(1), 71-100.

[7] Kyle, A. S. (1985). Continuous auctions and insider trading. Econometrica: Journal of the Econometric Society, 1315-1335.

[8] O'Hara, M. (1998). Market microstructure theory. John Wiley & Sons.

[9] Madhavan, A. (2000). Market microstructure: A survey. Journal of financial markets, 3(3), 205-258.

[10] Lux, T., & Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. Nature, 397(6719), 498-500.

[11] LeBaron, B. (2006). Agent-based computational finance. Handbook of computational economics, 2, 1187-1233.

[12] Tesfatsion, L., & Judd, K. L. (Eds.). (2006). Handbook of computational economics: agent-based computational economics (Vol. 2). Elsevier.

[13] Chen, S. H., & Yeh, C. H. (2001). Evolving traders and the business school with genetic programming: A new agent-based artificial stock market architecture. Journal of Economic Dynamics and Control, 25(3-4), 363-393.

[14] Patel, J., Shah, S., Thakkar, P., & Kotecha, K. (2015). Predicting stock and stock price index movement using trend deterministic data preparation and machine learning techniques. Expert systems with applications, 42(1), 259-268.

[15] Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European journal of operational research, 270(2), 654-669.

[16] Bao, W., Yue, J., & Rao, Y. (2017). A deep learning framework for financial time series using stacked autoencoders and long-short-term memory. PloS one, 12(7), e0180944.

[17] Heaton, J. B., Polson, N. G., & Witte, J. H. (2016). Deep learning in finance. arXiv preprint arXiv:1602.06561.

[18] Moody, J., & Saffell, M. (2001). Learning to trade via direct reinforcement. IEEE transactions on neural Networks, 12(4), 875-889.

[19] Nevmyvaka, Y., Feng, Y., & Kearns, M. (2006, June). Reinforcement learning for optimized trade execution. In Proceedings of the 23rd International Conference on Machine Learning (pp. 673-680).

[20] Jiang, Z., Xu, D., & Liang, J. (2017). A deep reinforcement learning framework for the financial portfolio management problem. arXiv preprint arXiv:1706.10059.