Multi-Agent Deep Reinforcement Learning for High-Frequency Multi-Market Making

Extended Abstract

Pankaj Kumar
Copenhagen Business School
Copenhagen, Denmark
pk.mpp@cbs.dk

ABSTRACT

High-frequency multi-market making is a liquidity-providing strategy that exercises cross-market latency arbitrage in order to simultaneously post multiple bids and asks in a fragmented market for a security or co- related securities, while maintaining a relatively low net position. By exploiting price discrepancies between markets, the strategy earns profit from the bid-ask spread for every trade against the risk of inventory, liquidity and adverse selection. We develop a multi-market simulation framework built over empirically verified heterogeneous agents, with a realistic market design and matching engine. We use it to design high-frequency market making agents based on deep attention recurrent Q-network architecture a with spatial and temporal attention module, to efficiently capture the non-linear features of the order book. We train heterogeneous market making agents, trading in the presence of other agents, with a simulation framework that employs independent Q-learning in a multi-agent deep reinforcement learning setting. We demonstrate the effectiveness of our agents in relation to traditional deep architecture and benchmark strategies using Deep Hawkes processes. We investigate the effect of latency and different market ecology on the market quality.

KEYWORDS

Multi-Agent Deep Reinforcement Learning; High-Frequency Market Making; Multi Agent-Based Model

ACM Reference Format:
Pankaj Kumar. 2023. Multi-Agent Deep Reinforcement Learning for High-Frequency Multi-Market Making: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

The proliferation of technological innovations in the financial market has led to a significant transformation in which conventional, human- driven floor trading is giving way to fragmented electronic exchanges. Furthermore, consequential changes in the microstructure of exchanges have contributed to exponential growth in the use of high-frequency trading (HFT, or high-frequency traders). Among the generalised trading strategies frequently utilised by HFT firms, high-frequency market making accounts for more than 60% of overall HFT volume [27]. The strategy exploits latency to earn profit in the form of the bid-ask spread, liquidity rebates and mid-quote changes. The rapidly growing literature following the Flash Crash of 6 May 2010 has polarised the discussion on high-frequency market making's effect on market quality, liquidity and price efficiency [1, 2, 12, 17, 18, 23]. However, there has been relatively little focus on the inherent characteristics of market making – namely multi-market effect, latency arbitrage and agents' profitability.

Agent-based models are acclaimed framework for studying market making problems in a single or multiple market setting [3, 4, 6, 13, 20, 21, 31–34]. They provide a framework within which to connect the behaviour of market makers to fragmented or single markets. This framework is then used to study the effect of market design (e.g. fragmented market, market integration, etc.) on market microstructure (e.g. liquidity, efficiency etc.). Deep reinforcement learning (DRL) approaches are a contemporary paradigm within which to investigate market making strategies [7, 9, 10, 19, 22, 29, 30]. The success of DRL led to DRL being integrated into multi-agent systems (MADRL). The agents trading in fragmented exchanges often have incomplete and noisy realization of the state of the market following partial observability. To address the earlier issue, Deep Recurrent Q-Network (DRQN) [15], Deep Attention Recurrent Q-Network (DARQN) [28], deep distributed recurrent Q-network (DDRQN), curriculum learning [11], deep recurrent policy inference Q-network (DRPIQN) [16], and Bayesian action decoder (BAD)[24] was introduced. Another widely used approach, independent Q-learning (IQL), which entails multi-agent fingerprints [8], was proposed in MADRL.

To the best of our knowledge, this paper is the first to study heterogeneous high-frequency market making strategies across fragmented markets in multi-agent deep reinforcement learning frameworks. High-frequency market making agents exploiting latency arbitrage strategies are trained using DARQN with prioritised experience replay [26] and double q-learning [14]. The training of multiple heterogeneous agents is implemented by IQL with multi-agent fingerprints to mitigate non-stationarity effects and effectively incorporate experience replay [8]. The development, training, validation and testing is done on a realistic simulation framework built over fragmented market design, latency enabled interface kernels, matching engine and the Financial Information eXchange (FIX) protocol [25]. We evaluate price, liquidity, volatility and the agents' profitability across markets from simulated order book event streams, by removing one agent at a time from the market ecology. The investigation contributes to the growing literature on interaction between agents trading on single or multiple markets. This is the first step toward capturing the essence of the trade web in the equity market, i.e. who feeds whom – the financial

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

equivalent of the food web. We also evaluate the effect of latency on the agent's profitability and market quality.

2 EXPERIMENTS AND RESULTS

The high-frequency market makers (DARQN, DRQN, DQN) are trained against benchmark high-frequency market makers (DHP), High Frequency Traders (HFT), Opportunistic Traders (OT), Fundamental Buyers/Sellers (FB/FS) and Small Traders (ST). Figure 1a, 1b and 1c illustrate the agents' performance. Next, We analyse the impact of latency on the agents' profitability and market quality, using our simulation framework, based on a two-market setup that features refined agent ecology. We reduce the latency of the original experimental setup by 10%, 20%, 30% and 40% to investigate changes in market quality and agents' profitability. For each latency configuration, we take only the test data of 500 trading days without changing the agents' ecology. Figure 1d, 1e and 1f illustrate the performance of the agents in different latency settings. Figure 1g, 1h and 1i shows market quality over a multiple latency configuration. *Depth(X)* [5], a measure of liquidity, varies significantly when latency is reduced.

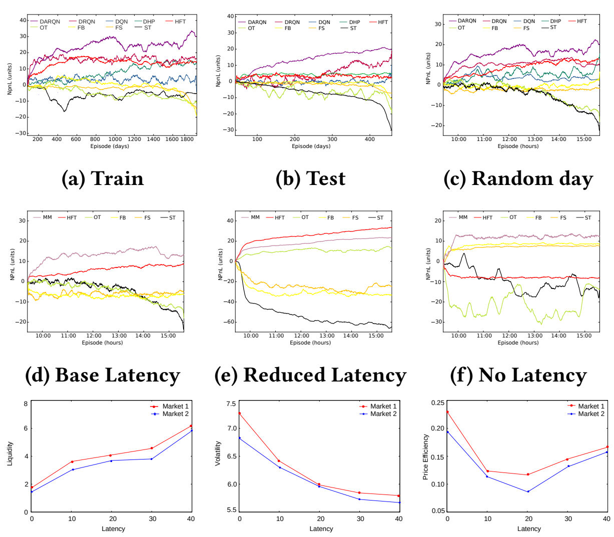

(a) Train (b) Test (c) Random day

(d) Base Latency (e) Reduced Latency (f) No Latency

(g) Liquidity (h) Volatility (i) Price Efficiency

Figure 1: Agents' performance and market quality for different latency configuration.

To investigate the trading profitability of different agents interacting with each other in our simulation framework, we knock out groups of agents from the base market ecology, one at a time. The analysis is based on the analogy of the "food web", i.e. who eats whom in an ecological community. Here, we might call it a "trade web", who feeds whom in a financial market. Figure 2 uses a box plot to show the agents' variance performance in different market ecologies. To understand market quality in presence of different market ecologies, we calculate *Depth(X)*, realised volatility and auto-correlation in mid-quote returns as market quality measures, averaged over 500 test days. Table 1 demonstrates liquidity, volatility and price efficiency in different market taxonomies.

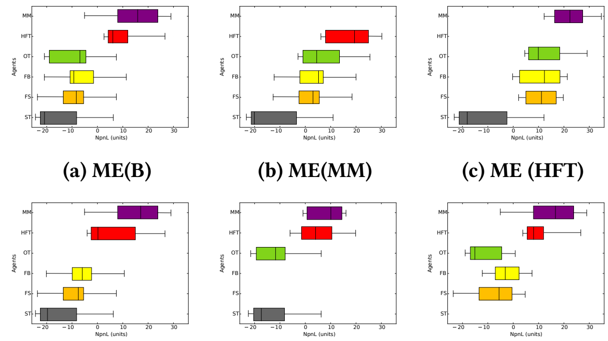

(a) ME(B) (b) ME(MM) (c) ME (HFT)

(d) ME (OT) (e) ME (FT) (f) ME (ST)

Figure 2: Agents' performance on a random day from the test set for different market ecology configurations. ME (B) represents the base market ecology, comprising MM, HFT, OT, FB, FS and ST. Groups of agents are knocked out one at a time to create different market ecologies. For example, ME (MM) represents a market ecology in which MM has been knocked out. The agents shown are median agents, and performance is reported in 10^5 unit currency.

Table 1: Market quality for different agent ecology configurations. *Depth(X)* represents liquidity, realised volatility represents volatility, and auto-correlation in mid-quote returns represents price efficiency. The values are reported as Market 1 (Market 2) in the unit specified.

	ME(B)	ME (MM)	ME (HFT)	ME (OT)	ME (FT)	ME (ST)
Liquidity	3.9(1.9)	2.6(1.4)	2.2(1.2)	3.3(1.5)	3.4(1.5)	3.9(1.9)
Volatility	4.2(3.7)	5.1(4.6)	7.1(6.2)	5.4(6.6)	5.3(7.6)	4.2(3.7)
Price Efficiency	0.13(0.12)	0.17(0.07)	0.16(0.10)	0.40(0.35)	0.70(0.62)	0.17(0.13)

3 CONCLUSIONS

This paper is the first to adapt the MADRL framework to multi-market making using a realistic simulation framework in fragmented markets. We model high-frequency maker making agents that exploit latency arbitrage using DARQN network architecture with spatial and temporal attention modules. The high-frequency market making agents interact with the simulation framework in an empirically verified market ecology and are trained using IQL with multi-agent fingerprints. Due to efficient learning ability, the DARQN agents outperform the other agents, in particular DRQN, DQN and DHP. We contribute to the growing literature on the effect of latency on the market quality. Our analysis finds that a reduction in latency improves liquidity and price efficiency, and reduces volatility. The investigation into the interaction of algorithmic trading strategies suggests that HFT and MM feed on orders from OT, FT and ST, the presence of which is presence is essential for improving market quality. However, when they compete among themselves, it has a detrimental effect on the market.

REFERENCES

[1] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. 2014. High-Frequency Trading and Price Discovery. *The Review of Financial Studies* 27, 8 (06 2014), 2267–2306.

[2] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. 2014. *Market Integration and High Frequency Intermediation.* Technical Report. Investment Industry Regulatory Organization of Canada (IIROC).

[3] Charles Collver. 2017. *An application of agent-based modeling to market structure policy: the case of the U.S. Tick Size Pilot Program and market maker profitability.* U.S. Securities and Exchange Commission. https://www.sec.gov/marketstructure/research/increasing-the-mpi-combined.pdf

[4] Sanmay Das. 2005. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance* 5, 2 (2005), 169–180.

[5] Hans Degryse, Frank de Jong, and Vincent van Kervel. 2015. The Impact of Dark Trading and Visible Fragmentation on Market Quality. *Review of Finance* 19, 4 (2015), 1587–1622.

[6] M Duffin and J Cartlidge. 2018. Agent-Based Model Exploration of Latency Arbitrage in Fragmented Financial Markets. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI).* 2312–2320.

[7] Marcus Elwin. 2019. *Simulating market maker behavior using Deep Reinforcement Learning to understand market microstructure.* Master's thesis. KTH Royal Institute of Technology, Stockholm.

[8] Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip H. S. Torr, Pushmeet Kohli, and Shimon Whiteson. 2017. Stabilising Experience Replay for Deep Multi-Agent Reinforcement Learning. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70* (Sydney, NSW, Australia) *(ICML'17).* JMLR.org, 1146–1155.

[9] Sumitra Ganesh, Nelson Vadori, Mengda Xu, Hua Zeng, Prashant Reddy, and Manuela Veloso. 2019. Reinforcement Learning for Market Making in a Multi-agent Dealer Market. arXiv:1911.05892

[10] Olivier Gueant and Iuliia Manziuk. 2019. Deep Reinforcement Learning for Market Making in Corporate Bonds: Beating the Curse of Dimensionality. *Applied Mathematical Finance* 26, 5 (2019), 387–452.

[11] Jayesh K. Gupta, Maxim Egorov, and Mykel Kochenderfer. 2017. Cooperative Multi-agent Control Using Deep Reinforcement Learning. In *Autonomous Agents and Multiagent Systems.* Springer International Publishing, Cham, 66–83.

[12] Björn Hagströmer and Lars Norden. 2013. The diversity of high-frequency traders. *Journal of Financial Markets* 16, 4 (2013), 741 – 770.

[13] Thomas A Hanson. 2016. High frequency traders in a simulated market. *Review of Accounting and Finance* 15, 3 (08 2016), 329–351.

[14] Hado van Hasselt, Arthur Guez, and David Silver. 2016. Deep Reinforcement Learning with Double Q-Learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence* (Phoenix, Arizona) *(AAAI'16).* AAAI Press, 2094–2100.

[15] Matthew Hausknecht and Peter Stone. 2015. Deep Recurrent Q-Learning for Partially Observable MDPs. In *AAAI Fall Symposium on Sequential Decision Making for Intelligent Agents (AAAI-SDMIA15)* (Arlington, Virginia, USA).

[16] Zhang-Wei Hong, Shih-Yang Su, Tzu-Yun Shann, Yi-Hsiang Chang, and Chun-Yi Lee. 2018. A Deep Policy Inference Q-Network for Multi-Agent Systems. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) *(AAMAS '18).* International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1388–1396.

[17] Andrei Kirilenko, Albert Kyle, Mehrdad Samadi, and Tugkan Tuzun. 2017. The Flash Crash: High Frequency Trading in an Electronic Market. *The Journal of Finance* 72, 3 (2017), 967–998.

[18] Robert Korajczyk and Dermot Murphy. 2018. High-Frequency Market Making to Large Institutional Trades. *The Review of Financial Studies* 32, 3 (2018), 1034–1067.

[19] Pankaj Kumar. 2020. Deep Reinforcement Learning for Market Making. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Auckland, New Zealand) *(AAMAS '20).* International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1892–1895.

[20] Sandrine Jacob Leal, Mauro Napoletano, Andrea Roventini, and Giorgio Fagiolo. 2016. Rock around the clock: An agent-based model of low- and high-frequency trading. *Journal of Evolutionary Economics volume* 26 (08 2016), 49–76.

[21] Zhuoshu Li and Sanmay Das. 2016. An Agent-Based Model of Competition Between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs?. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls (Eds.). ACM, 50–58.

[22] Iwao Maeda, David deGraw, Michiharu Kitano, Hiroyasu Matsushima, Hiroki Sakaji, Kiyoshi Izumi, and Atsuo Kato. 2020. Deep Reinforcement Learning in Agent Based Financial Market Simulation. *Journal of Risk and Financial Management* 13, 4 (2020), 1–17.

[23] Albert Menkveld. 2013. High frequency trading and the new market makers. *Journal of Financial Markets* 16 (11 2013), 712–740.

[24] T Nguyen, D Nguyen, and S Nahavandi. 2020. Deep Reinforcement Learning for Multiagent Systems: A Review of Challenges, Solutions, and Applications. *IEEE Transactions on Cybernetics* 50, 9 (2020), 3826–3839.

[25] Consulting Oxera. 2009. *What are the benefits of the FIX Protocol? Standardising messaging protocols in the capital markets.* Technical Report. Prepared for FIX Protocol Limited.

[26] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. 2016. Prioritized Experience Replay. In *International Conference on Learning Representations.* Puerto Rico.

[27] SEC. 2014. *Equity market structure literature review part ii: High frequency trading.* U.S. Securities and Exchange Commission, Staff of the Division of Trading and Markets, 63–103. http://www.sec.gov/rules/concept/2010/34-61358

[28] I Sorokin, A Seleznev, M Pavlov, A Fedorov, and A Ignateva. 2015. Deep Attention Recurrent Q-Network. *ArXiv* abs/1512.01693 (2015).

[29] Thomas Spooner, John Fearnley, Rahul Savani, and Andreas Koukorinis. 2018. Market Making via Reinforcement Learning. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) *(AAMAS '18).* International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 434–442.

[30] Thomas Spooner and Rahul Savani. 2020. Robust Market Making via Adversarial Reinforcement Learning. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems* (Auckland, New Zealand) *(AAMAS '20).* International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2014–2016.

[31] Elaine Wah, Dylan Hurd, and Michael Wellman. 2015. Strategic Market Choice: Frequent Call Markets vs. Continuous Double Auctions for Fast and Slow Traders. *EAI Endorsed Transactions on Serious Games* 3, 10 (8 2015).

[32] Elaine Wah and Michael P. Wellman. 2013. Latency Arbitrage, Market Fragmentation, and Efficiency: A Two-Market Model. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce* (Philadelphia, Pennsylvania, USA) *(EC '13).* Association for Computing Machinery, New York, NY, USA, 855–872.

[33] Y Wang, H Zhou, and W Zeng. 2011. The agent-based simulation of inventory-based model's impact on market maker trading mechanism. In *The Fourth International Workshop on Advanced Computational Intelligence.* 601–604.

[34] Y Xiong, T Yamada, and T Terano. 2015. Comparison of different market making strategies for high frequency traders. In *2015 Winter Simulation Conference (WSC).* 324–335.