
On the Importance of Opponent Modeling in Auction Markets

Mahmoud Mahfouz^{*1,2}, Angelos Filos^{*1,3}, Cyrine Chtourou¹, Joshua Lockhart¹
Samuel Assefa¹, Manuela Veloso¹, Danilo Mandic², Tucker Balch¹
¹ JPMorgan Chase & Co., ² Imperial College London, ³ University of Oxford
mahmoud.a.mahfouz@jpmchase.com *

Abstract

The dynamics of financial markets are driven by the interactions between participants, as well as the trading mechanisms and regulatory frameworks that govern these interactions. Decision makers would rather not ignore the impact of other participants on these dynamics, and should employ tools and models that take this into account. To this end, we demonstrate the efficacy of applying *opponent-modeling* in a number of simulated market settings. While our simulations are simplified representations of actual market dynamics, they provide an idealized “playground” in which our techniques can be demonstrated and tested. We present this work with the aim that our techniques could be refined and, with some effort, scaled up to the full complexity of real-world market scenarios. We hope that the results presented encourage practitioners to adopt opponent-modeling methods and apply them on live systems, in order to enable not only reactive but also proactive decisions to be made.

1 Introduction

The study of multi-agent systems (MAS) plays an important role in different scientific disciplines, including, but not limited to biology [19], evolutionary psychology [14], economics [16], and political science [27]. MAS has also been studied in computational finance to develop trading agents [35] and their associated decision making strategies, taking into account the interactions with the other market participants. In MAS, the agents attempt to maximize a utility function which is typically dependent on the behavior and actions of the other agents in the environment. This suggests the need for agent strategies to be reactive and adaptive to the other agents in the environment.

Game theory [29, 38] provides a mathematical formalism for describing the complex behaviors that emerge from the different agent interactions and suggests tools to develop successful strategies. Nonetheless, strong assumptions are traditionally made about having access to reliable models describing the system dynamics, which is invalid even in the simplest of setups. For example, in financial markets designing a strategy that takes into account the impact of its decisions, is as difficult as solving the original problem itself, and heuristic approaches [18] are not adaptive hence their quality degrades as the dynamics drift over time.

In contrast, learning from data [29] is an attractive framework under which assumptions can be relaxed and flexible models can be developed and calibrated to real-world data. In particular, recent advances in machine learning [17, 26] allow for black-box, task-agnostic, large-scale modeling. However, in all these settings, the presence of multiple learning agents renders the training problem non-stationary and often leads to unstable dynamics or undesirable final results [4, 15].

*Equal contribution

Opponent modeling [11] addresses this problem, attempting to *decouple the sources of non-stationarity* by modeling separately the other learning agents from the stationary components of the environment.

In this paper we explore the applications of opponent modeling techniques in single and continuous double auction markets [35], as a proxy for real financial markets. We also discuss potential research topics that address key challenges and opportunities of interest to the financial community.

2 Background and Related Work

Opponent modeling refers to a set of methods to construct models of the other agents in the same environment with the goal of predicting various properties of interest such as the agents’ actions, types, beliefs and goals [1]. Different approaches are studied in the literature including policy reconstruction [8, 33], type-based reasoning [5, 2], classification [34], plan recognition [10], recursive reasoning [12], graphical models [22] and group modeling [30].

The financial markets can be thought of as a multi-agent system composed of multiple agents whose different utility functions (*e.g.* maximizing profit or reducing risk exposure) are highly dependent on the behavior and actions of the other market participants. For instance, Yang et al. [39, 40] implicitly perform opponent modeling in a continuous double auction setup by using Inverse Reinforcement Learning (IRL) to characterize the different trading strategies in a limit order book. In their approach, they model the trades placed in the order book as a Markov Decision Process (MDP) and cluster the reward space to differentiate between high frequency, opportunistic and market making strategies.

3 Opponent Modeling in Auction Markets

We formulate decision-making in a multi-agent system as that of finding an optimal policy in a Markov game (MG) [25]. An MG for N agents is defined by a set of states \mathcal{S} , N sets of actions $\mathcal{A}_1, \dots, \mathcal{A}_N$ for each agent. The next state is produced according to the state transition function $\mathcal{T} : \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_N \mapsto \mathcal{S}$, and at each state, agent i receives reward $r_i : \mathcal{S} \mapsto \mathbb{R}$. Each agent i aims to maximize its own total expected return $R_i = \sum_{t=0}^T \gamma^t r_i^t$ where γ is a discount factor and T is the time horizon. Note that the transition function depends on *all* agents’ actions, hence each agent has an incentive to model and adapt to the other agents actions too, since they contribute to its own expected return. We distinguish between two types of policies for solving an MG:

(1) *implicit-modeling*: the agents implicitly take into account the impact of the other agents via the global state, learning a policy $\pi_i^{(IM)} : \mathcal{S} \times \mathcal{A}_i \mapsto [0, 1]$;

(2) *opponent-modeling*: each agent explicitly models the policies of other agents, learning a model $\hat{\pi}_{-i}$ and choose actions according to $\pi_i^{(OM)} : \mathcal{S} \times \mathcal{A}_i \times \hat{\pi}_{-i} \mapsto [0, 1]$, where the subscript $-i$ refers to all agents but i .

Note that the implicit-modeling approach is treating the MG as a single-agent Markov decision process (MDP) [6], but there are no guarantees for the existence of a stationary optimal policy since the transition dynamics depend on uncontrollable variables (*i.e.*, all agents but i).

An attractive framework for learning policies for sequential decision-making from data and without access to reliable models is model-free reinforcement learning [31] which is also suitable for solving MGs. For agent i , we assume parametric policy π_{θ_i} , which we train using the simple REINFORCE algorithm [37, 32], directly adjusting the parameters θ_i in order to maximize the objective $J(\theta_i) = \mathbb{E}_{s \sim p^\pi, a \sim \pi_{\theta_i}} [R_i]$ by taking steps in the direction of $\nabla_{\theta_i} J(\theta_i)$.

In this section we present two examples of auction games, reformulated as MGs and demonstrate the superiority of opponent modeling over opponent-agnostic approaches. Despite the simplicity and abstraction of these examples, links and relevance are drawn with real financial market applications in Section 4. We build on the analysis discussed in [35] for auctions and present two representative examples:

3.1 First-Price Sealed Bid Iterated Auctions

In a first-price sealed bid eponymous iterated auction [35] (also known as a blind auction), all bidding agents simultaneously submit private bids (that is, no bidding agent knows of the other agent’s bids) which are revealed at the end of each round. The highest bidder wins, and pays the price they submitted. This process is repeated T times.

Explicitly, consider N agents, each with a *fixed* and private evaluation $v_i \in [0, 1]$. The i -th agent’s policy π_{θ_i} returns a bid $b_i \in [0, 1]$ and the reward received is $v_i - b_i$ if the agent i won the auction (i.e., submitted the highest bid) or 0 otherwise. A rational agent does not submit bids larger than its evaluation and tries to maximize the margin $v_i - b_i$ subject to winning the auction. Note that the game is *stateless*.

We assume that $N - 2$ agents sample from a latent truncated Gaussian distributions with different means $\{\mu_i\}_{i=1}^{N-2}$ and variances $\{\sigma_i^2\}_{i=1}^{N-2}$, maximum value at $\{v_i\}_{i=1}^{N-2}$ and minimum value of 0.

The remaining two agents are adaptive and trained using the REINFORCE algorithm [37, 32]. The first agent (IM) with policy $\pi_{\theta_{IM}}$, performs implicit modeling. We chose a truncated Gaussian distribution for our model, hence $\theta_{IM} = \{\mu_{IM}, \sigma_{IM}^2\}$ with maximum and minimum values v_{IM} and 0, respectively. The second adaptive agent performs opponent modeling and learns a model of the other agents using maximum likelihood estimation for $\hat{\pi}_{-OM}$ which is used to find the best response $\pi_{\theta_{OM}}$.

Figure 1 summarizes the results from a simulation with $N = 10$, $T = 50$ and $v_{IM} = v_{OM} > v_i$, for all $i \in \{1, \dots, N - 2\}$. In this experiment, the agent performing opponent modeling wins 58.2% of the auctions demonstrating the value of our approach. The agent performing implicit modeling and the $N - 2$ random agents win 34.5% and 7.3%, respectively.

3.2 Continuous Double Auctions

Previously, we assumed that direct opponent modeling was possible, since the actions of each agent were announced *after* an auction, hence explicit models $\hat{\pi}_{-i}$ could be built on the historic actions. In many cases this information is not available even a posteriori, such as in limit order books of exchanges.

A limit order book (LOB) is used by exchanges to match buyers and sellers of a given security according to price-time priority rules [7]. The LOB represents a snapshot of the supply and demand for a given security at a given time. While exchanges maintain historical records of the orders submitted to the limit order book, the identity and the type of the trading participants is not typically revealed. This makes the problem very difficult to approach in a supervised manner using only observations of the order book and the orders placed in the market.

We create labeled *synthetic* data using an agent-based limit order book simulator (ABIDES) [9] and formulate our approach to opponent modeling in a limit order book as a supervised learning problem of identifying (classifying) agent *archetypes* explicitly programmed in the simulator, a task which can be cast in the auxiliary tasks framework [23, 21] for semi-supervised reinforcement learning. Using a classification approach for opponent modeling [34] allows for predicting properties of the other agents as opposed to predicting the future actions of the agents, which can be useful in a range of market scenarios discussed in our work.

The trading archetypes considered for generating our synthetic data are: (1) *market makers (MM)*, (2) *liquidity consumers (LC)*, (3) *mean reversion traders (ME)* and (4) *momentum traders (MO)*. For details about the trading agents implemented, please refer to the appendix.

Synthetic Data: The synthetic market is created in an open source agent-based limit order book simulator, ABIDES [9]. The simulation is run for a trading day between 09:30 AM and 4:00 PM using the specified pool of agents acting against a market replay agent submitting the historical orders




Figure 1: Percentage of auctions won per agent type, where the opponent-modeling (OM) agent dominates both the implicit-modeling (IM) and random policies.

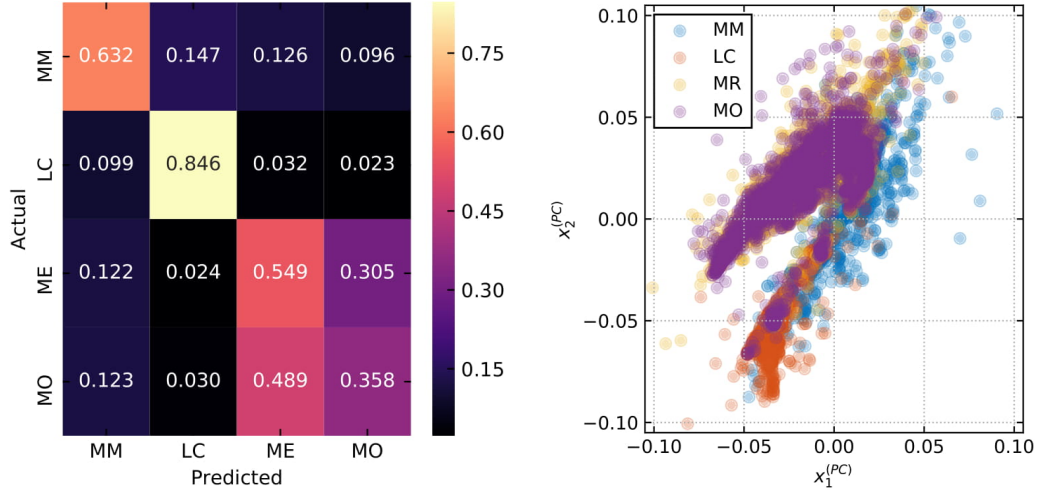


Figure 2: Trading agent archetype classification in a limit order book: model results

for IBM on the 28th of June 2019. The market created consisted of 67 market makers, 500 liquidity consumers, 500 mean reversion traders and 500 momentum agents. An issue arises regarding the different time scales that each agent archetype operates on. For example, some archetypes place substantially more orders than others, leading to the data set we generate having imbalance between the classes. We produce a balanced dataset by simply down-sampling the over-represented classes. The dataset used to train the model consists of the limit order book snapshots representing the state-space, individual agent order details representing the action-space and the agent archetype as the label. The state-space comprises the two best bid and ask prices and sizes in the limit order book prior to the agent’s order. We also compute other features including the mid price $M(t)$, the last 5 observed midprices by the agent, the spread (the difference between the best ask price $A(t)$ and best bid price $B(t)$), volume imbalance (ratio of the best ask size to the best bid size) and the order flow imbalance [13], an imbalance metric tracking the net order flow in the bid and ask queues. The action-space contains the order details including the direction of the trade, the price and size of the order, the relative size of the order with respect to the opposite queue and the difference between the order price and the mid price. The synthetic data set created consists of 20,000 state-action pairs (5000 state-action pairs per trading agent archetype).

Model: The classification model used is a simple feed-forward neural network that attempts to infer the agent archetype given the limit order book state and the agent’s order. The model consists of two hidden layers with 128 units each. The activation function used is ReLU with a 0.0004 learning rate and an ADAM optimizer [24]. The train-validation-test split used was 53.6 %, 13.4% and 33.4% of the state-action pairs respectively and the model was run for 1000 epochs.

Results: Using the synthetic data generated and the model discussed, we obtain an 88% classification accuracy in predicting the different agent archetypes: market maker (MM), liquidity consumer (LC), mean reversion trader (ME) and momentum trader (MO) in the test set. The confusion matrix for the multi-class classification model is shown in Figure 2a. While our model performs well on classifying the different archetypes, the performance is lower in differentiating between the momentum and mean reversion trading agents. This is expected given the myopic nature of the classifier and the training data used to train the model. In Figure 2b, we apply Principal Component Analysis on the feature maps of the penultimate layer of the neural network and project this onto 3 principal components to get an understanding of how the model learns to differentiate between the classes. It is clear that there is a high degree of overlap between the different classes when looking at the first two components which suggests that the problem is not easy to tackle even in the simple setup proposed.

4 Conclusion and Future Work

Other than the toy examples demonstrated in Section 3, methods of opponent modeling in financial multi-agent systems can be useful in a number of scenarios and from different perspectives, including but not limited to:

Regulation and Safety: Market manipulation strategies represent an element of price distortion and creation of artificial market conditions, which in turn impedes market efficiency, transparency and fairness. In the LOB context, spoofing agents, for instance, submit orders only to cancel them before their execution and then submit a buy or sell order to benefit from the change that their initial (fake) order caused. The ability to explicitly model the policies of such agents could help in the early detection of market abuse and avoiding adverse consequences like the flash crash that happened on May 6, 2010, causing a large and quickly reverting dip in the Dow Jones Industrial Average stocks. Moreover, *Level 6* information, consisting of LOB labeled records, are exclusively available to exchanges and regulators to verify that market participants are complying with rules and processes [28]. Having access to such a data set is an opportunity to make a more informed and dynamic utility modeling. Quantifying proportions of market makers, liquidity takers and buy side participants would be valuable in detecting signs of recession or unhealthy market conditions.

Decision Making and Automation: Market participants usually only have access to (at most) *Level 4* LOB data [28], which is *anonymized* order flow data with new, modified, and canceled limit orders placed at a given snapshot. Although this poses a challenge to the disambiguation of limit orders placed in the order book because it imposes a myopic approach as opposed to using a consistent timeseries of consecutive orders placed by the same agent, unlabeled data does not fall short of information and opportunities. Opponent modeling can be used for the detection of iceberg orders and other forms of hidden liquidity [20]. By studying the occurrence of such orders, *a posteriori*, in the LOB and the surrounding market conditions, we can attempt to model the utility behind such strategies. Ref. [20] studied the correlation between dark orders and market conditions and showed that market conditions reflected by the bid-ask spread, depth and price movements significantly impact the aggressiveness of ‘dark’ liquidity supply. If a trader can detect such dark orders through recognizing a pattern that iceberg order placing agents fail to conceal or through modeling the utility of such agents and the specific cases that make them submit such orders, they can estimate and take into account these orders’ effect on the LOB balance and predict a larger than observed spread.

Additionally, the Request for Quotes (RFQ) process is another setting where we could benefit from opponent modeling. RFQs are used to facilitate trading in illiquid markets *e.g.* fixed income exchange-traded funds (ETFs). In this setting, clients send quote requests to multiple dealers for a specific quantity of a security. Ultimately, the dealer with the most attractive offer wins the auction. In this single, reverse, and static auction setting, one buyer (the auctioneer/the client) is looking to buy a single security from multiple sellers (the dealers) who do not observe the offers submitted by the other dealers, and only receive feedback from the auctioneer once the latter’s decision for the winner of the auction is made. It is therefore important for dealers receiving RFQs from different clients to have a good estimate of the expected level at which the other dealers are willing to sell the security to avoid mis-evaluating the offer price. This can in theory be accomplished by having a good estimate of the fair value of the asset in question. However, it is difficult to estimate the fair value of the asset if the asset is illiquid. If a participant is able to infer the utilities, and by extension, the expected offer prices of each of the other dealers in the auction, they can narrow down their uncertainty about the fair value of the asset and avoid mispricing it. In RFQs, apart from modeling competitors, it can be extremely valuable to be able to model clients. In this case data is labeled, as sell side firms (banks) are aware of their RFQ counterparty (client). Performing efficient client tiering would allow to respond to client orders in an efficient manner, which is beneficial for both the client and the bank.

Finally, opponent modeling applied to financial markets, can be a step forward in the automation of decision making. Indeed, explicitly inferring the different market participant strategies would allow to obtain a tailored, well defined utility function that is inspired from high-performing market participants.

We have explored the application of opponent modeling approaches to the problem of building autonomous agents that operate within an auction setting. We have shown how opponent modeling can be applied in the single and double auction settings. In the former case, we showed that reinforcement learning agents that directly model their opponents in a single auction market outperform those that

do not. In the double auction setting, we considered the problem of identifying participants in a limit order book from the actions they perform in response to the current state of the order book. We showed that a supervised learning approach can be used to distinguish between different types of market participants in a controlled environment using synthetic data.

The simple experiments explored and the initial results obtained are encouraging evidence that opponent modeling is potentially feasible and useful in settings of interest to the different actors in the financial markets. Future work in this direction could explore more unsupervised or semi-supervised learning approaches especially in the case of the continuous double auction settings. However, the challenges anticipated following such approaches would be around labeling any clusters obtained which would still require domain knowledge. This suggests the need for both supervised and unsupervised learning approaches.

Disclaimer

This paper was prepared for information purposes by the AI Research Group of JPMorgan Chase & Co and its affiliates ("J.P. Morgan"), and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no explicit or implied representation and warranty and accepts no liability, for the completeness, accuracy or reliability of information, or the legal, compliance, financial, tax or accounting effects of matters contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

References

- [1] Stefano V Albrecht and Peter Stone. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258:66–95, 2018.
- [2] Stefano Vittorino Albrecht, Jacob William Crandall, and Subramanian Ramamoorthy. E-hba: Using action policies for expert advice and agent typification. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- [3] Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [4] David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. *arXiv preprint arXiv:1802.05642*, 2018.
- [5] Samuel Barrett, Peter Stone, Sarit Kraus, and Avi Rosenfeld. Teamwork with limited knowledge of teammates. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [6] Richard Bellman et al. The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6):503–515, 1954.
- [7] Hendrik Bessembinder. Price-time priority, order routing, and trade execution costs in nyse-listed stocks. *Order Routing, and Trade Execution Costs in Nyse-Listed Stocks (June 2001)*, 2001.
- [8] George W Brown. Iterative solution of games by fictitious play. *Activity analysis of production and allocation*, 13(1):374–376, 1951.
- [9] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity market simulation for ai research. *arXiv preprint arXiv:1904.12066*, 2019.
- [10] Sandra Carberry. Techniques for plan recognition. *User Modeling and User-Adapted Interaction*, 11(1-2):31–48, 2001.
- [11] David Carmel and Shaul Markovitch. Opponent modeling in multi-agent systems. In *International Joint Conference on Artificial Intelligence*, pages 40–52. Springer, 1995.

- [12] David Carmel and Shaul Markovitch. Incorporating opponent models into adversary search. In *AAAI/IAAI, Vol. 1*, pages 120–125, 1996.
- [13] Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88, 2014.
- [14] Leda Cosmides and John Tooby. Evolutionary psychology and the generation of culture, part ii: Case study: A computational theory of social exchange. *Ethology and sociobiology*, 10(1-3): 51–97, 1989.
- [15] Jakob Foerster, Richard Y Chen, Maruan Al-Shedivat, Shimon Whiteson, Pieter Abbeel, and Igor Mordatch. Learning with opponent-learning awareness. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 122–130. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- [16] James W Friedman. *Game theory with applications to economics*, volume 87. Oxford University Press New York, 1986.
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [18] Markus Gsell. Assessing the impact of algorithmic trading on markets: A simulation approach. Technical report, CFS Working Paper, 2008.
- [19] Peter Hammerstein and Reinhard Selten. Game theory and evolutionary biology. *Handbook of game theory with economic applications*, 2:929–993, 1994.
- [20] Nikolaus Hautsch and Ruihong Huang. On the dark side of the market: identifying and analyzing hidden order placements. Available at SSRN 2004231, 2012.
- [21] Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. Agent modeling as auxiliary task for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 15, pages 31–37, 2019.
- [22] Ronald A Howard and James E Matheson. Influence diagrams. *Decision Analysis*, 2(3):127–143, 2005.
- [23] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. *arXiv preprint arXiv:1611.05397*, 2016.
- [24] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [25] Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pages 157–163. Elsevier, 1994.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- [27] James D Morrow. *Game theory for political scientists*. Princeton University Press, 1994.
- [28] Mark Paddrik, Roy Hayes, William Scherer, and Peter Beling. Effects of limit order book information level on market stability metrics. *Journal of Economic Interaction and Coordination*, 12(2):221–247, 2017.
- [29] Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.
- [30] Peter Stone and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 110(2): 241–273, 1999.

- [31] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [32] Richard S Sutton, David A McAllester, Satinder P Singh, and Yishay Mansour. Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pages 1057–1063, 2000.
- [33] Manuela M Veloso. *Planning and learning by analogical reasoning*, volume 886. Springer Science & Business Media, 1994.
- [34] Ben G Weber and Michael Mateas. A data mining approach to strategy prediction. In *2009 IEEE Symposium on Computational Intelligence and Games*, pages 140–147. IEEE, 2009.
- [35] Michael P Wellman. Trading agents. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(3):1–107, 2011.
- [36] J Welles Wilder. *New concepts in technical trading systems*. Trend Research, 1978.
- [37] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- [38] Michael Wooldridge. *An introduction to multiagent systems*. John Wiley & Sons, 2009.
- [39] Steve Yang, Mark Paddrik, Roy Hayes, Andrew Todd, Andrei Kirilenko, Peter Beling, and William Scherer. Behavior based learning in identifying high frequency trading strategies. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFER)*, pages 1–8. IEEE, 2012.
- [40] Steve Y Yang, Qifeng Qiao, Peter A Beling, William T Scherer, and Andrei A Kirilenko. Gaussian process-based algorithmic trading strategy identification. *Quantitative Finance*, 15(10):1683–1703, 2015.

Appendix: Trading Agents

Market Makers: These are market participants who aim to provide liquidity in the LOB by placing orders on both sides of the order book. That is, given the market's current mid price $M(t)$, a market maker will place orders at prices $B(t) = M(t) - p_b(t)$ and $A(t) + p_a(t)$ where $p_a(t)$ and $p_b(t)$ are chosen according to some strategy. Market makers operate under the assumption that there is enough trading activity for the instrument that they are operating in, and that the aggregate buying and selling of that instrument will net them a profit proportional to the difference in the prices they offer $A(t) - B(t)$. This difference is referred to as the spread offered by the market maker. In our experiments we utilize two basic market making strategies: The first places two orders (one buy and one sell order) at the top of the book (*i. e.* the highest bid price and lowest ask price in the order book respectively) at each time step. The volumes of the orders are chosen at random following a uniform distribution within fixed bounds v_{\min} and v_{\max} . The second strategy takes a similar approach but places orders deeper into the book at prices $1, \dots, N$ levels above and below $M(t)$ for N chosen randomly from a uniform distribution within fixed bounds N_{\min} and N_{\max} . In addition, we include an implementation of the Avellaneda-Stoikov strategy [3], which dynamically adjusts its bid and ask prices as a function of volatility and the amount of inventory held by the agent.

Liquidity Consumers: In our setup, liquidity consumers represent directional traders placing market orders to buy/sell a given quantity at the best available price. Each agent is configured to place a buy or sell order with the direction and order size chosen at random from a uniform distribution.

Momentum Traders: The momentum agents base their trading decision on observed price trend signals. Our implementation compares the 20 past mid price observations with the 50 past observations and places a buy order of random size, if the former exceeds the latter and a sell order otherwise.

Mean Reversion Traders: These represent agents placing orders in the market following mean reversion signals. In our setup, we have two groups of agents following these type of signals. The first looks at the divergence from an exponentially weighted moving average signal and places buy/sell orders if the price deviates above/below a certain number of standard deviations away from the average. The second group utilizes a similar approach using the Relative Strength Index (RSI) signal [36] instead. The RSI is a technical indicator that is used to detect and identify mean reversion. It is defined in terms of the exponential moving average, and the average of up and down price moves over a set period. Although usually the RSI is used on close prices, we will be using a micro-structure variant of RSI where we will look at the time series of mid prices at the tick level. Traditionally, [36], when RSI is greater than the 70 level, the stock is considered to be overbought, and when the indicator is lower than the 30 level, the stock is considered to be oversold. Our agent therefore places a sell order when the stock is overbought, and a buy order when it is oversold and makes no action when the stock is neutral.