Deep Hawkes Process for High-Frequency Market Making

Extended Abstract

Pankaj Kumar
Jheronimus Academy of Data Science
’s-Hertogenbosch, Netherlands
p.kumar1@tue.nl

ABSTRACT

High-frequency market making is a liquidity-providing trading strategy that simultaneously generates many bids and asks for a security at ultra-low latency while maintaining a relatively neutral position. The strategy makes a profit from the bid-ask spread for every buy and sell transaction, against the risk of adverse selection, uncertain execution and inventory risk. We design realistic simulations of limit order markets and develop a high-frequency market making strategy in which agents process order book information to post the optimal price, order type and execution time. By introducing the Deep Hawkes process to the high-frequency market making strategy, we allow a feedback loop to be created between order arrival and the state of the limit order book, together with self- and cross-excitation effects. Our high-frequency market making strategy accounts for the cancellation of orders that influence order queue position, profitability, bid-ask spread and the value of the order. The experimental results show that our trading agent outperforms the baseline strategy, which uses a probability density estimate of the fundamental price. We investigate the effect of cancellations on market quality and the agent’s profitability. We validate how closely the simulation framework approximates reality by reproducing stylised facts from the empirical analysis of the simulated order book data.

KEYWORDS

High-Frequency Market Making; Multi Agent-Based Model; Deep Hawkes Process; Stacked Denoising Autoencoder

ACM Reference Format:
Pankaj Kumar. 2024. Deep Hawkes Process for High-Frequency Market Making: Extended Abstract. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 3 pages.

1 INTRODUCTION

Technological innovations and regulatory initiatives have shifted the financial market from traditional exchange floors to electronic exchanges, dominated by algorithmic trading. High-frequency trading (HFT) accounts for 85% of equity market trading, with a focus on rapid order execution and complex strategies [6, 9]. The Flash Crash of 2010 reignited discussions on HFT’s role in a fragmented market.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

Research on market making spans finance, agent-based modeling, and artificial intelligence [2–5, 11]. Traditional stochastic optimal control models and information-based models have limitations, leading to exploration of alternative approaches. In recent years, deep learning, particularly the Deep Long-Short Term Memory (DLSTM) architecture, has gained prominence in high-frequency finance [10]. Stacking Denoising Autoencoders (SDAEs) address issues like random weight initialization. This paper introduces the Deep Hawkes process (DHP), combining SDAE with DLSTM. This paper is the first to incorporate DHP into high-frequency market making, allowing feedback loops and enhancing predictive power. DLSTM-SDAE outperforms the Neural Hawkes process (NHP) in predicting order types and times.

The paper develops a scalable multi-asset simulation framework based on realistic market architecture. The framework introduces a feedback loop between order arrival and the state of the order book using DHP in the high-frequency market making setting. However, it has limitations in accounting for basket events and external factors. The paper explores predictive and trading performance, including the impact of cancellation on various market parameters [8].

2 RESULTS

2.1 Predictive Performance

Figure 1 illustrates high-frequency market making agents’ predictive performance on reconstructed order book data, trained as per [7], at nanosecond and millisecond resolutions.

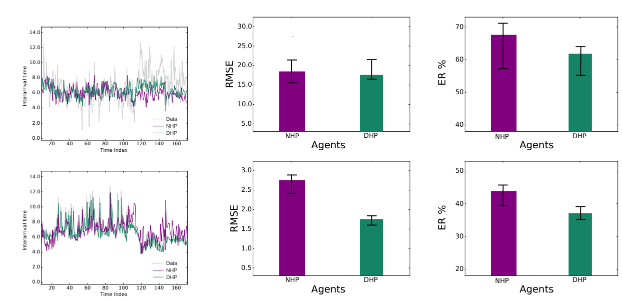

(a) Temporal **(b) Temporal Error** **(c) Event Error**

Figure 1: Evaluation of high-frequency market making agents’ performance in predicting order book events and time at nanosecond (upper) and millisecond (lower) resolution. Error bars show standard deviation over 10 experiments.

2.2 Trading Performance

In Figure 1, DHMM outperforms NHMM in predicting order type and time, crucial for high-frequency market making. The DLSTM-SDAE architecture enables agents to learn hidden representations from noisy order book data, enhancing profitability. DHMM exhibits faster convergence, as seen in Figure 2. The PMM baseline strategy relies on a probability density estimate linked to the fundamental price's jump process. Initially profitable, PMM's performance diminishes over time as DHMM and NHMM master precise order placement. Long-term, PMM might excel with extended holding periods, warranting investigation into its performance under different probability density estimate conditions on microstructure feature joint distributions.

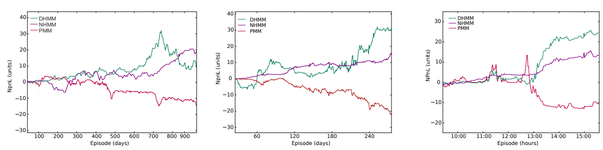

(a) Train (b) Test (c) Random day

Figure 2: Trading agents performance with DHMM, NHMM and PMM while training, testing and random day.

2.3 Order Cancellation Effect

We examine profit, bid-ask spread, and order queue position distributions by removing the cancellation mechanism in the simulation framework. The order's intrinsic value concerning queue position is estimated using [8]'s model. The agent's queue position, reflecting the orders ahead at a specific price, impacts execution speed, fill rate, latency, and adverse selection cost. Queue position is estimated by reconstructing the limit order book from the simulated data feed. For empirical calibration, we utilize parameters from [8], encompassing order size distribution, trade arrival rate (TAR), average trade size (ATS), trade size in the standard lot (TSS), cancellation arrival rate (CAR), average cancellation size (ACS), price jump arrival rate (PJR), average jump size (AJS), market impact (MI), and average queue size (AQS). Table 1 details the estimated parameters for simulated data without cancellation mechanisms (Simulated NC), with cancellation mechanisms (Simulated WC), and an average over 21 days (Simulated AV).

Table 1: Estimated parameters for simulated orderbook data.

Data	TAR (/min)	ATS (shares)	TSS (shares)	CAR (/min)	ACS (shares)	PJR (/min)	AJS (ticks)	MI	AQS (shares)
Simulated NC	2.53	3467	7664	92.72	5061	1.26	0.32	10.91	16416
Simulated AV	2.04	4037	6901	82.21	4022	1.01	0.46	8.76	23554
Simulated WC	5.26	6329	8083	43.71	1107	3.75	2.06	11.92	40191
Simulated AV	4.07	5463	9147	40.31	1560	3.01	2.06	13.02	46815

Table 1 reveals that the absence of cancellation mechanisms significantly increases the average queue size for high-frequency market makers. The reduction in cancellation rate amplifies the queue size, impacting the market maker's profitability, bid-ask spread, and market impact. The order's value, considering queue position, bid-ask spread, and agent's profit, aligns with the patterns shown in Figure 3. Wider bid-ask spreads, as illustrated in Figure 3d when limit order cancellations are restricted, adversely affect profitability (Figure 3e) compared to scenarios with cancellations (Figure 3a,3b). According to the model, the value of an unfilled order is zero. Figure 3f demonstrates that an increase in queue length decreases the probability of execution and, consequently, the order's value. The value becomes flat as the queue length becomes extremely large. Our results corroborate findings by [1] in exploring determinants of order cancellations. While causal relationships are challenging to infer in artificially created scenarios, this approach lays the groundwork for future investigations using order-level data.

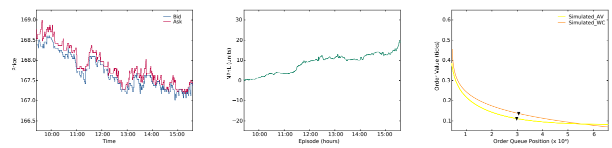

(a) BAD WC (b) DHMMP WC (c) VOPQ WC

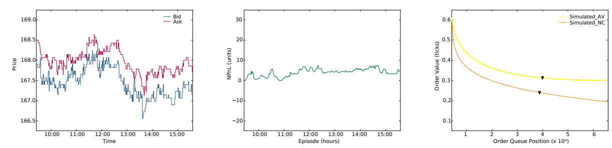

(d) BAD NC (e) DHMMP NC (f) VOPQ NC

Figure 3: Effect of Limit Order Cancellations on Market Dynamics. The top row (WC) illustrates the impact when market maker's agents can cancel limit orders, while the bottom row (NC) shows a scenario with no cancellations. BAD represents intraday bid-ask distribution, DHMMP displays profit distribution of DHMM throughout the trading day, and VOPQ denotes the value of orders relative to queue position. The black triangle signifies the average queue length on a specific trading day.

3 CONCLUSIONS

We have formulated a high-frequency market-making strategy that integrates feedback loops, self-modulating multivariate Hawkes processes, and DLSTM-SDAE within our realistic simulation framework. Applying this strategy to nanosecond-resolution reconstructed order book data reveals a subpar performance in predicting the next order type and timestamp. However, with millisecond-resolution data, the strategy surpasses NHP in prediction tasks and benchmark market making strategies. The extension of DHP in a market making setting further improves performance, affirming empirical claims regarding the impact of cancellations on order size determinants. Our model points to potential research avenues, including advanced pre-processing and learning algorithms for noise filtering, exploration of intensity-free approaches for Hawkes processes, and integration of DHP into reinforcement learning frameworks for optimal policy learning. Future work includes extending the model to deep reinforcement learning, direct extraction of trading algorithm parameters from order book data, and delving into causality within diverse research domains.

REFERENCES

[1] Petter Dahlström, Björn Hagströmer, and Lars L. Norden. 2024. Determinants of Limit Order Cancellations. *Financial Review* 59, 1 (2024), 181–201.

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

[3] Lawrence Glosten and Paul Milgrom. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14, 1 (1985), 71 – 100.

[4] Thomas Ho and Hans Stoll. 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9, 1 (Mar 1981), 47 – 73.

[5] 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.

[6] Pankaj Kumar. 2023. Multi-Agent Deep Reinforcement Learning for High-Frequency Multi-Market Making. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems* (London, United Kingdom) *(AAMAS '23)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2409–2411.

[7] Hongyuan Mei and Jason Eisner. 2017. The Neural Hawkes Process: A Neurally Self-Modulating Multivariate Point Process. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) *(NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 6757–6767.

[8] Ciamac C. Moallemi and Kai Yuan. 2017. *A model for queue position valuation in a limit order book*. Technical Report. Columbia Business School Research Paper No. 17-70.

[9] Purba Mukerji, Christine Chung, Timothy Walsh, and Bo Xiong. 2019. The Impact of Algorithmic Trading in a Simulated Asset Market. *Journal of Risk and Financial Management* 12 (04 2019), 68.

[10] Alaa Sagheer and Mostafa Kotb. 2019. Time series forecasting of petroleum production using deep LSTM recurrent networks. *Neurocomputing* 323 (2019), 203 – 213.

[11] 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.