Get Real: Realism Metrics for Robust Limit Order Book Market Simulations

Svitlana Vyetrenko[1], David Byrd[1,3], Nick Petosa[3], Mahmoud Mahfouz[1,2], Danial Dervovic[1], Manuela Veloso[1], and Tucker Hybinette Balch[1]

[1]J. P. Morgan AI Research, New York, NY, USA *
[2]Department of Electrical and Electronic Engineering, Imperial College, London, UK
[3]School of Interactive Computing, College of Computing, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

Machine learning (especially reinforcement learning) methods for trading are increasingly reliant on simulation for agent training and testing. Furthermore, simulation is important for validation of hand-coded trading strategies and for testing hypotheses about market structure. A challenge, however, concerns the robustness of policies validated in simulation because the simulations lack fidelity. In fact, researchers have shown that many market simulation approaches fail to reproduce statistics and stylized facts seen in real markets. As a step towards addressing this we surveyed the literature to collect a set of reference metrics and applied them to real market data and simulation output. Our paper provides a comprehensive catalog of these metrics including mathematical formulations where appropriate. Our results show that there are still significant discrepancies between simulated markets and real ones. However, this work serves as a benchmark against which we can measure future improvement.

1 Background and related work

1.1 Motivation

Most professional investors, hedge funds, investment institutions and banks need robust means of testing trading strategies in simulation before "going live" with funds at risk. A key reason for this is to gain assurance that the strategy is likely to be effective. *Alpha* strategies aim to profit from price movements, while *execution* strategies are intended to complete large volume orders while minimizing transaction costs. For instance, a pension fund may have concluded that it should reduce its holdings in a particular stock and therefore trigger a sell order for that asset. If this order were sent to an exchange as a market sell order, the price would likely fall significantly and provide the seller a lower average price than they would hope. In order to reduce transaction costs, it is a common practice to design execution strategies so that price impact is minimized by distributing a larger order as a set of smaller orders over time [1].

Significant research effort is aimed at applying Reinforcement Learning to a variety of trading problems in which the learners are trained in simulation: A reinforcement learning market-maker was presented in [2]; a reinforcement learning approach to algorithmic execution was introduced in

---

[3]; deep hedging a portfolio of derivatives (including over-the-counter derivatives) in the presence of market friction was considered in [4]; LSTM representations for an RL trading agent was given in [5].

In these financial trading problems, the statistical properties of the environment and how the market responds to trading activity is often unknown and difficult to model. In such cases, repeating the process a number of times in a simulated environment allows us to eliminate the need of knowing transition probabilities explicitly, and an optimal policy can be learned from the gained simulated experience requiring realistic market simulation tools.

In real-time algorithmic trading, the actions of any given agent incurs response from other market participants. In simulation, autonomous agents can choose to place orders at any time and market response to them will not be reflected in historical data. Therefore, simple market replay of historical data is not sufficient for back testing or strategy construction. Interactive agent-based simulation (IABS) can potentially simulate interaction between individual market participants [6]. In such simulators, price arises from incentives of fully autonomous agents each of whom act rationally in order to maximize their profits. These principles reflect how real markets operate; the challenge is, however, to find realistic agent configurations and prescribe agent behavior in such a way that their actions produce synthetic time series whose statistical properties resemble real markets.

1.2 Stylized Facts of Limit Order Book (LOB) Behavior

Later sections of this paper rely on the reader's understanding of the mechanisms by which electronic markets operate, so we briefly review them here. Public exchanges such as NASDAQ and NYSE facilitate the buying and selling of assets by accepting and satisfying buy and sell orders from multiple market participants. The exchange maintains an order book data structure for each asset traded. The LOB represents a snapshot of the supply and demand for the asset at a given time. It is an electronic record of all the outstanding buy and sell limit orders organized by price levels. A matching engine, such as first-in-first-out, is used to match incoming buy and sell order interest [7].

Order types are further distinguished between limit orders and market orders. A limit order specifies a price that should not be exceeded in the case of a buy order, or should not be gone below in the case of a sell order. Hence, a limit order queues a resting order in the LOB at the side of the book of the market participant. A market order indicates that the trader is willing to accept the best price available immediately. A diagram illustrating LOB structure is provided in Figure 1.
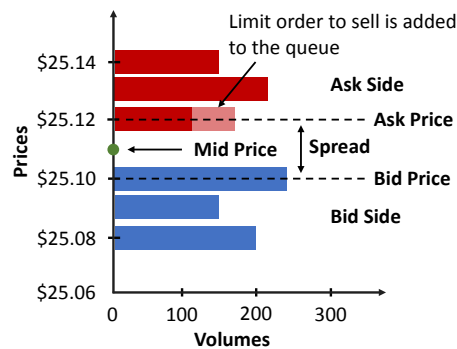

Figure 1: Visualization of the LOB structure.

Many of the metrics we present below are derived from the observation of the LOB over time. Properties of LOB behavior that are repeated across a wide range of instruments, markets, and time periods are referred to as *stylized facts* [8].

Evaluating the statistical properties of simulated asset returns, order volumes, order arrival times, order cancellations, etc and comparing them to those generated from real historical data allows us to infer the level of fidelity of a simulation. The question of whether stylized facts originate from traders' behavior, or if they are a natural consequence of order book markets, has been widely discussed in the literature. If some stylized facts can be derived from markets populated only by zero intelligence (ZI) agents which make decisions without the knowledge of market microstructure, then these facts must

originate from the mechanism that is governing the markets and not from strategic agent behavior. For example, [9] showed that a market simulation that consists only of ZI agents that place limit and market orders independently at random is able to reproduce price and spread dynamics as well as market impact. Ability of ZI agents to reproduce fat tails and long range dependence was shown in [10], however, the authors emphasized that in order to reproduce correlated order phenomena such as volatility clustering, strategic agent behavior may be required.

In this paper, we provide a survey of several groups of LOB stylized facts across multiple asset classes that lead to realism metrics with respect to empirical distributions (defined in Section 2). We compare two simulator configurations: one that contains zero intelligence (ZI) agents only and another that combines ZI agents with minimal strategic behavior agents. We find that the more diverse agent configuration behaves more similarly to real markets; however, we conclude that for more robust simulation of collective market phenomena online learning adaptive agents might be necessary. [11, 10].

1.3 Related work

IABS methods allow us to study phenomena that emerge as a consequence of multiple participant interactions and are difficult to model otherwise. Examples of such complex phenomena include both the routine market microstructure events such as market response to an individual participant's trading [12]; and rare events such as flash crashes [13, 14, 15] as well as extreme market shocks.

Wellman helped establish an empirical approach to the study of markets using simulated multi-agent systems [16]. In this approach, referred to as Empirical Game Theoretic Analysis (EGTA), many tens of thousands of simulations are run during which the strategies of agents are adjusted until the system reaches a Nash equilibrium. It is only after the system reaches equilibrium that its statistical properties are evaluated. We regard this as a desirable approach, but we have not yet implemented the methods necessary for discovering equilibria in our experimental platform.

One notable example of multi-agent simulator use success is the NASDAQ tick size experiment where NASDAQ researchers experimentally demonstrated that under some agent scenarios reducing tick size would actually lead to increased spreads (an undesirable property) and would negatively impact price discovery [17, 18]. These findings are counter-intuitive and illustrate importance of multi-agent simulations for market policy research since one would expect that smaller tick size results in tighter spreads. Some explanations of these observations include the possibility of the existence and proliferation of parasitic strategies that can take advantage of better prices to make the market more jittery and volatile. These findings are particularly important since NASDAQ is interested in finding policies that have not yet been discovered and used by market participants in order to prevent potential market manipulations.

In real-time trading, injecting orders to the market induces other market participant activity that typically drives prices away from the agent. This activity is known as market impact [1, 19]. Presence of market impact in real time implies that a realistic trading strategy simulation should include deviation from historical data. In literature, it is common to make an assumption of negligible market impact given the size of agent orders is small and sufficient amount of time is allowed between consecutive trades [20]. A simple two-agent simulated market environment that consists of an algorithmic trading agent and the rest of the market is presented in [21]. It is a partially synthetic data model that allows to deviate from historical data only at times when agent places market orders which are known to cause highest market impact. This model is, however, only suited for small order placement, and is unable to capture more complex dynamics of transient price impact [22, 23].

While modeling the market as an interplay of multiple agents seems a natural approach to mimic real market collective emergent behavior, justifying the realism of such approach for validating new trading strategies is difficult. Agent modeling typically relies on common sense hand-crafted rules (e.g., [10]), which can be difficult to calibrate as historical data labeled with details about each individual constituent agent behavior is typically not available for public use. Several calibration approaches— e.g. error minimization to find parameters for the asset pricing model with heterogeneous beliefs [24] and using Bayesian parameter estimation techniques in the simulated context—have been introduced [25]. When individual agent- or execution strategy-specific data is available to the researcher, it can be used for the simulator calibration (e.g., [21, 26]).

One can view the multi-agent LOB environment as a non-cooperative game in which every agent pursues their own goal and there is no communication among agents [27]. From a game-theoretic

perspective, a simulated environment is more realistic when it reaches a Nash equilibrium, where every player's parameters are configured so that each can do no better by unilaterally changing its strategy. Agents that learn to maximize their long term rewards by reinforcement from empirical equilibrium environments have been discussed in [11].

Other approaches to IABS realism can include inverse learning agents' rewards from the market [26]; generating synthetic LOB data using GANs [28]; incorporating feedback from real-time trading into the simulation [29] and building adaptive agents that are governed by the evolutionary principles and can learn from experience [30, 31].

2 Realism metrics

We now focus on and review a series of metrics of LOB behavior found in the literature. One way to establish IABS realism is to ensure that simulated LOB time series mimic the stylized facts derived from real market histories. Below we review several groups of such stylized facts across multiple asset classes [26, 32, 9, 10].

2.1 Notation and definitions

For simplicity of presentation, we introduce some notation and definitions that will be used throughout this paper. At time t, let b_t be the best bid price, and let a_t be the best ask price. We define mid-price as $m_t = \frac{a_t + b_t}{2}$. Choose time scale Δt. Given a time scale Δt, which can range from milliseconds to months, the log return (or simply return) at scale Δt is defined as $r_{t,\Delta t} = \ln m_{t+\Delta t} - \ln m_t$. Let $\sigma_{\Delta t}$ be return volatility which can be calculated as standard deviation of price returns.

Let x be size of a new order placed into LOB. Let T be a lifetime of an order until it is fully executed or canceled. We denote by $b_t - \Delta$ the price of a new buy limit order, and $a_t + \Delta$ the price of a new sell limit order. Notice that Δ can be negative. Let V_a and V_b be the volumes available at the best bid and ask price. Slice LOB price and volume time series in small non-overlapping time intervals. For each time interval τ, let μ_{V_τ} be the average traded volume and $\sigma_{\tau,\Delta t}$ be the return volatility over τ. Furthermore, let $P(.)$ denote probability density function of a given quantity.

2.2 Stylized facts about asset return distributions

Multiple stylized facts about price return distributions were studied in [8] for equity markets as well as in [33] for foreign exchange and rates markets.

- **Absence of autocorrelations** Linear autocorrelations $\text{corr}(r_{t+\tau,\Delta t}, r_{t,\Delta t})$ of asset returns over periods τ longer than 20 minutes are insignificant.

- **Heavy tails and aggregational normality** The distribution of daily asset price returns shows fat tails; however, as one increases the period of time Δt over which these returns are calculated, asset returns show lower tails. One way to quantify deviation from normal distribution is to calculate its kurtosis.

- **Intermittency** At any micro or macro time scale, asset price returns must display high degree of volatility.

- **Volatility clustering** High-volatility events tend to cluster in time. A quantity used to measure volatility clustering is the autocorrelation function of the squared returns $\text{corr}(r_{t+\tau,\Delta t}^2, r_{t,\Delta t}^2)$. Empirical studies using returns from various equities indicate that this autocorrelation function remains significantly positive over several days, which indicate periods of high volatility clustering [8].

- **Long range dependence** If one looks at autocorrelation function of absolute returns as a function of time lag $f(\tau) = \text{corr}(|r_{t+\tau,\Delta t}|, |r_{t,\Delta t}|)$, it is empirically shown that it decays according to the power law distribution $f(\tau) \sim \tau^{-\beta}$ with exponent $\beta \in [0.2, 0.4]$ [8].

- **Gain/loss asymmetry** Gain/loss asymmetry is prevalent for equity price returns as stocks lose value faster than they grow [8]. However, this trend is not as pronounced for foreign exchange and rates products. Skewness is a metric that can be used to quantify the asymmetry of probability distribution about its mean.

4

- **Volume/volatility positive correlation** Volume and volatility are positively correlated. Linear regression relationship $\mu_{V_\tau} \sim \alpha + \beta \sigma_{\tau, \Delta t}$ can be derived from the data [34].
- **Returns/volatility negative correlation** Asset returns/volatility are negatively correlated.
- **Asymmetric causal information flow** Coarse-scaled volatility predicts fine-scaled volatility better than fine-scaled volatility predicts coarse scaled-volatility,

2.3 Stylized facts about volumes and order flow

- **Order book volumes** Volumes at best bid V_b (and respectively volumes at best ask V_a) are distributed according to Gamma distribution for $\gamma \leq 1$ [35]:
$$P(V_b) \sim \exp^{-V_b} V_b^{-1+\gamma}.$$
- **Order sizes** Order sizes are power-law distributed [7]. For instance,[36] show examples when limit order sizes are distributed as $P(x) \sim x^{-(1+\mu)}$ with exponent $1 + \mu \approx 2$ and market order sizes are distributed as $P(x) \sim x^{-(1+\mu)}$ with exponent $1 + \mu \approx 2.3 - 2.7$. Orders tend to have round number of shares (i.e. multiples of 10, 100, etc. are more common than neighboring sizes); in general, power-law distribution fit is product-specific.
- **Number of orders in a fixed time window** Number of orders in a fixed time window can be approximated by gamma or lognormal distributions [36].
- **Order inter-arrival times** In the literature, LOB order inter-arrival times are suggested to be fit into exponential [28], lognormal, and Weibull distributions [36].
- **New order prices** Prices at which new limit orders are placed, are power-law distributed around bid-ask [36]. Specifically, $P(\Delta) \sim \Delta^{-(1+\mu)}$ with $1 + \mu \approx 1.6$ [35].
- **Cancellation time, time-to-first-fill and time-to-execution** Lifetimes of both cancelled and executed limit orders are power-law distributed, $P(T) \sim T^{-(1+\mu)}$ with $1 + \mu$ ranging between 1.3 and 1.6 for both canceled and executed limit orders [36]. Since an order can require multiple fills to be completed, one must distinguish between time-to-first-fill and time-to-completion statistics. Generalized gamma distribution with accelerated failure time can be used to model time-to-first-fill and time-to-completion distributions [37].
- **Time correlation of order flow** Individual agent's order placement decisions depend on other agents' actions [10].

2.4 Stylized facts about non-stationary patterns

- **Intraday volume patterns** LOB volumes are known to exhibit strong intraday patterns. For instance, historical foreign exchange trading volumes can be approximated by fifth-degree polynomial "U-shaped" regional sessions that correspond to New York, London, and Tokyo trading [38]. Similarly, in most equity markets, volumes are highest in the beginning of trading day, followed by a period of lower activity, and then spike again at the end of the trading day [7]. Note that making a transformation from physical time to tick (or transaction) time may help adjusting for intraday non-stationarity [39].
- **Seasonal volume patterns** Some assets, especially those consumer demand for which is seasonal (e.g., electricity futures), display strong seasonal volume patterns.
- **Intraday sensitivity to macro economic events/holidays** Due to product sensitivity to macro factors, volume spikes are known to occur in foreign exchange and rates markets during economic announcements. Equities trading is also sensitive to economic events [40]. Additionally, lower trading volumes are observed on holidays throughout all asset classes.
- **Intraday volume/spread negative correlation** Lower spreads are typically observed during periods of higher trading volumes.

2.5 Stylized facts about order market impact

Market impact of order placement is a expected to grow as a function of order volume. For each time interval τ, define $V_{\text{buy}, \tau}$ and $V_{\text{ask}, \tau}$ to be buy and sell order volumes in τ respectively. Define participation of volume in τ as
$$P_\tau = \frac{|V_{\text{buy}, \tau} - V_{\text{ask}, \tau}|}{V_{\text{buy}, \tau} + V_{\text{ask}, \tau}}.$$

Note that $0 \leq P_\tau \leq 1$. Also define Δm_τ to be the observable mid-price move in τ. Discretize the range for P_τ into bins $B_i, i = 1, \ldots, N$ such that $B_i = \{\tau : \frac{i-1}{N} \leq P_\tau \leq \frac{i}{N}\}$. For each B_i, define

$$M_i = \frac{1}{|B_i|} \sum_{\tau \in B_i} \Delta m_\tau \quad \text{and} \quad P_i = \frac{1}{|B_i|} \sum_{\tau \in B_i} \Delta P_\tau$$

to be the average price move and average participation of volume in bins with similar volume participation. One can then fit the relationship of the form $M_i \sim \alpha P_i^\beta$ through the data [19, 9, 41].

2.6 Stylized facts about cross asset correlations

When simulating multiple assets, cross asset correlation properties must hold. For instance, equity index and its major constituents must show high degree of correlation [8]. For futures, for example, asset price moves across term structure are highly correlated and exhibits consistent patterns uncovered by the PCA (e.g., [33]). It is also worth noting that extreme returns (e.g., 99-percentile returns that occur during financial crisis) across various stocks or asset classes can be extremely correlated while their average returns are not [8].

3 Experiments with a multi-agent simulation

3.1 ABIDES - Agent-Based Interactive Discrete Event Simulation environment

In order to evaluate the ability of a given agent configuration to reproduce stylized facts about the market, we employ ABIDES, an agent-based interactive discrete event simulation environment [42]. ABIDES provides a selection of background agent types (such as agent types described in Section 3.2), a NASDAQ-like exchange agent which lists any number of securities for trade against a LOB with price-then-FIFO matching rules, and a simulation kernel which manages the flow of time and handles all inter-agent communication. Trading agents may not inspect the state of the exchange directly, but must direct realistic messages to request order book depth, obtain last trade prices, or place or cancel limit orders through the kernel, which imposes delays for computational effort and communication latency. Time proceeds in nanoseconds and all securities are priced in cents. The ABIDES source code is freely available under a BSD-style license at `https://github.com/abides-sim/abides`.

3.2 Background agent types

In order to conduct experiments, we specify one ZI and three minimal strategic agent types (market maker, momentum and heuristic belief learning agents). Detailed explanations and examples for the ZI and HBL agents and the fundamental process which drives them are available in [43].

Zero intelligence agents: Originally introduced in [44], ZI agent class includes a variety of agents that do not base their trading decisions on the knowledge of LOB microstructure. Similarly to [45], ZI agents in our implementation are enhanced with knowledge of noisy observation of exogenous 'true' value of stock which represents agent's understanding of the outside world (eg. earnings reports, macro events, immediate trading demand). ZI agents in our experiment arrive to the market according to a Poisson process and do not use LOB microstructure signals for trading decisions.

Heuristic belief learning (HBL) agent: HBL agents base their decisions on a limited-length historical snapshot of the order stream, which they use to maximize expected surplus using a heuristic estimation of the probability that a given limit price will successfully transact in the market [46], [47]. Our implementation matches that of [32].

Market maker agent: The market maker agent acts as a liquidity provider by placing orders on both sides of LOB with a constant arrival rate of 10 seconds. The agent starts by cancelling any existing orders. It then queries the current spread to determine the prices of the buy and sell orders to be submitted. The agent is configured to place orders on top N levels on both sides of LOB with the size split determined by the number of price levels it quotes and the order volume chosen uniformly at random within fixed bounds. Our implementation is similar to that of [48, 49].

Momentum agents: The momentum agents base their trading decision on observed price trends. 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.

3.3 Agent configurations

In this paper we present two background agent configurations, one that is composed exclusively of ZI agents (*sparse_zi_100*—contains 100 ZI agents), and one that has both ZI and strategic agents (*rmsc01*—contains 1 market maker agent, 50 ZI agents, 25 HBL agents and 24 momentum trading agents). Note that in both configurations, all agents are permitted to reenter the market. To assess the realism of *sparse_zi_100* and *rmsc01* ABIDES configurations, we look at stylized facts about asset returns and order flow described in Section 2.

4 Experimental results

4.1 Stylized facts about asset return distributions

To derive historical asset return distributions, we analyze minutely intraday log returns of 50 randomly sampled U.S. exchange-traded equities for each trading day of 2011. The set of equities is resampled each trading day and is drawn uniformly across all stocks from all exchanges.

- **Heavy tails and aggregation normality.** We confirm experimentally that both *rmsc01* and *sparse_zi_100* show heavy tails and aggregational normality, with *rmsc01* being closer to historical data (see Figure 2).
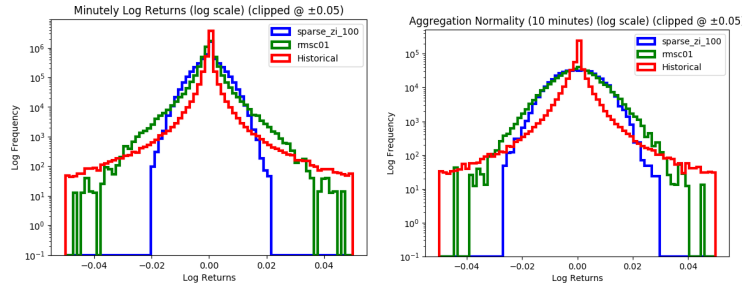

Figure 2: One-minute (left) and ten-minute (right) log return distributions.

- **Absence of autocorrelations** Figure 3 (*left*) shows the correlation coefficient distributions for lag 1 autocorrelation of log returns over 30 minute intervals. Historical market data exhibits a notable spike of 0 correlation, in line with previous discussion in Section 2. Both the *sparse_zi_100* and *rmsc01* agent configurations fail to capture this stylized fact. When comparing the overall distribution shape, the *rmsc01* configuration resembles the shape of historical autocorrelation distribution closer.

- **Volatility clustering.** The average autocorrelation of square returns decays for both historical and simulated data as time lag increases (see Figure 3 (*center*)).

- **Volume/volatility correlation.** Neither simulated distribution of correlation coefficients between volume and volatility was as skewed or high-variance as that of historical data (see Figure 3 (*right*)).

4.2 Stylized facts about volumes and order flow

To derive historical distributions, we consider order book historical data for JPM stock traded on the NASDAQ exchange for each trading day of June 2019 from 9:30 am to 4:30 pm.

- **Number of orders in a fixed time window.** Figure 4 (*left*) shows limit order volume distribution in a five-minute window for the simulated vs. the historical data. We find that gamma distribution produces a good fit for these curves.

- **Intraday volume patterns.** Quadratic curves have been fitted to this data to demonstrate the "U-shaped" pattern of historical intraday volumes (see Figure 4 (*right*)). The simulation data for *rmsc01* shows the reverse of expected activity: namely, reduced trading at the open and close of the market.
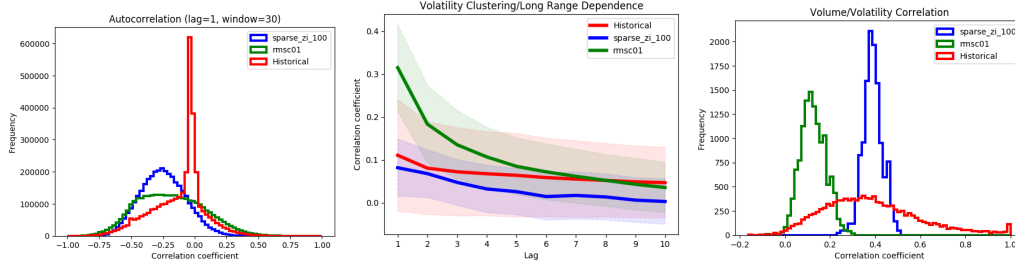
Figure 3: (*left*) Distributions of return autocorrelation. (*center*) Average autocorrelation of square returns as a function of time lag. (*right*) Volume/volatility correlation distributions.

- **Order interarrival times.** Figure 4 (*right*) shows distribution of interarrival times for simulated data as well as the historical data. Historical data fits Weibull distribution rather well. The distribution for *rmsc01* is dominated by values close to zero resulting in a poor fit to the Weibull distribution.
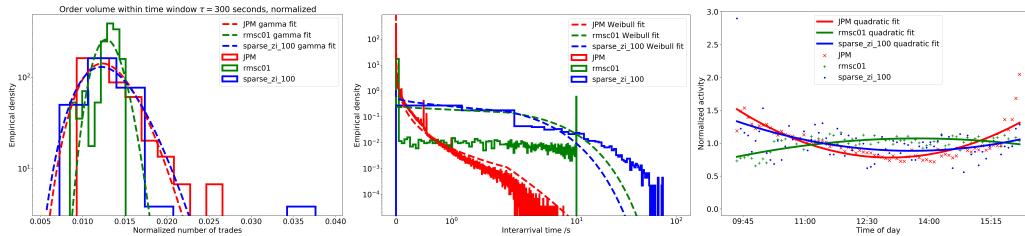

Figure 4: (*left*) Number of orders in a five-minute window. (*center*) Limit order interarrival times. (*right*) Intraday volume profiles.

5   Conclusion and Discussion

In this paper, we provided a catalog of known stylized facts regarding LOB microstructure behavior with respect to market realism. We chose ABIDES environment [42] as a simulation platform and evaluated two experimental configurations of agent types—one exclusively of ZI agents and one with ZI agents and other agents with minimal strategic behavior.

We observed that the configuration with a more diverse agent population leads to statistics that more closely mimic real markets, we acknowledge that there is much room for improvement. In particular, time correlation of order flow and intraday volume patterns are not well reproduced. While additional constraints might be needed to produce stylized facts about intraday patterns from agents' incentives (e.g., require agents to close all positions by the end of the trading day), correlated order behaviors, especially herding or clustering behaviors, require adaptation of one agent's behavior in response to other agents' actions and will possibly require introduction of online learning agents [10]. For example, [30] conducted comparisons of non learning and learning agents and concluded that agents capable of learning and adaption to other agent flows are able to replicate stylized facts about long range dependence and correlation between volume and volatility better. Specifically, during periods of high volumes, when more agents are trading in response to others' behavior, higher trading activity keeps volume queues available at best bid or ask levels relatively short; hence, LOB layers move more frequently and, as a result, prices are more volatile. Moreover, in real markets, rational agents evolve over time by learning to expand effective and cull ineffective trading strategies [31]. Hence, we believe that enhancing autonomous LOB agents with ability to learn from experience will be a step towards making simulated environments more robust.

6   Acknowledgements

References

[1] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 1999.

[2] T. Spooner, J. Fearnley, R. Savani, and A. Koukorinis. Market making via reinforcement learning. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, pages 434–442, Stockholm, Sweden, 2018.

[3] Y. Nevmyvaka, Y. Feng, and M. Kearns. Reinforcement learning for optimized trade execution. In *Proceedings of the 23rd International Conference on Machine Learning*, Pittsburg, PA, 2006.

[4] H Buehler, L Gonon, J Teichmann, and B Wood. Deep hedging. *arXiv:1802.03042*, 2018.

[5] David Lu. Agent inspired trading using recurrent reinforcement learning and lstm neural networks. *arXiv:1707.07338*, 07 2017.

[6] C Macal and M North. Tutorial on agent-based modeling and simulation. *Journal of Simulation*, 2010.

[7] Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier, and Martin Gould. *Trades, quotes and prices: financial markets under the microscope*. Cambridge University Press, Cambridge, 2018.

[8] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. 2001.

[9] J Farmer, Paolo Patelli, and Ilija Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences of the United States of America*, 102:2254–9, 03 2005.

[10] Imon Palit, Steve Phelps, and Wing Lon Ng. Can a zero-intelligence plus model explain the stylized facts of financial time series data? pages 653–660, 06 2012.

[11] L. Julian Schvartzman and Michael P. Wellman. Stronger cda strategies through empirical game-theoretic analysis and reinforcement learning. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, 2009.

[12] Wei Cui and Anthony Brabazon. An agent-based modeling approach to study price impact. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering and Economics*, pages 1–8, 03 2012.

[13] Sandrine Leal and Mauro Napoletano. Market stability vs. market resilience: Regulatory policies experiments in an agent based model with low- and high-frequency trading. *Journal of Economic Behavior and Organization*, 12 2016.

[14] Mark Paddrik, Roy Hayes, Andrew Todd, Steve Yang, Peter Beling, and William Scherer. An agent based model of the e-mini s&p 500 applied to flash crash analysis. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFEr)*, pages 1–8. IEEE, 2012.

[15] CFTC. Findings regarding the market events of may 6, 2010. *Report of the staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.*, 2010.

[16] Michael P Wellman. Methods for empirical game-theoretic analysis. In *AAAI*, pages 1552–1556, 2006.

[17] Vincent Darley and Alexander V. Outkin. *Nasdaq Market Simulation: Insights on a Major Market from the Science of Complex Adaptive Systems*. 2007.

[18] Eric Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 06 2002.

[19] Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. Direct estimation of equity market impact. *RISK*, 18, 04 2005.

[20] Thomas Spooner, John Fearnley, Rahul Savani, and Andreas Koukorinis. Market making via reinforcement learning. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, 2018.

[21] Svitlana Vyetrenko and Shaojie Xu. Risk-sensitive compact decision trees for autonomous execution in presence of simulated market response. In *ICML 2019 Workshop on AI in Finance*, 06 2019.

[22] Jean-Philippe Bouchaud, Yuval Gefen, Marc Potters, and Matthieu Wyart. Fluctuations and response in financial markets: the subtle nature of random price changes. *Quantitative Finance*, 07 2003.

[23] Jim Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 2010.

[24] Gabriele Tedeschi, Maria Recchioni, and Mauro Gallegati. A calibration procedure for analyzing stock price dynamics in an agent-based framework. *http://ssrn.com/abstract=2330739*, 09 2013.

[25] Jakob Grazzini, Matteo Richiardi, and Efthymios Tsionas. Bayesian estimation of agent-based models. *Journal of Economic Dynamics and Control*, 11 2015.

[26] S. Yang, M. Paddrik, R. Hayes, A. Todd, A. Kirilenko, P. Beling, and W. Scherer. Behavior based learning in identifying high frequency trading strategies. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering Economics (CIFEr)*, 03 2012.

[27] Junling Hu and Michael P. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning*, ICML '98, 1998.

[28] Junyi Li, Xintong Wang, Yaoyang Lin, Arunesh Sinha, and Michael P. Wellman. Generating realistic stock market order streams, 2019.

[29] N. Ruiz, S. Schulter, and M. Chandraker. Learning to simulate. In *Proceeding of the International Conference on Learning Representations*, New Orleans, LA, 2019.

[30] Blake LeBaron and Ryuichi Yamamoto. Long-memory in an order-driven market. 2007.

[31] Andrew Lo. Reconciling efficient markets with behavioral finance: The adaptive markets hypothesis. *Journal of Investment Consulting*, 7, 01 2004.

[32] Xintong Wang and Michael P Wellman. Spoofing the limit order book: An agent-based model. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 651–659. International Foundation for Autonomous Agents and Multiagent Systems, 2017.

[33] Giuseppe Ballocchi, Michel M. Dacorogna, Carl M. Hopman, Ulrich A. Muller, and Richard B. Olsen. The intraday multivariate structure of the Eurofutures markets. *Journal of Empirical Finance*, 6(5):479–513, December 1999.

[34] Olivier Brandouy, Angelo Corelli, Iryna Veryzhenko, and Roger Waldeck. A re-examination of the "zero is enough" hypothesis in the emergence of financial stylized facts. *Journal of Economic Interaction and Coordination*, 7:223–248, 10 2012.

[35] Jean-Philippe Bouchaud, Marc Mezard, and Marc Potters. Statistical properties of stock order books: Empirical results and models. *Quantitative Finance*, 2:301, 04 2002.

[36] Frédéric Abergel, Marouane Anane, Anirban Chakraborti, Aymen Jedidi, and Ioane Muni Toke. *Limit order books*. Cambridge University Press, 2016.

[37] Andrew Lo, A Craig MacKinlay, and June Zhang. Econometric models of limit-order executions. *Journal of Financial Economics*, 65, 03 2001.

[38] Michael M. Dacorogna, Ulrich Alfons Müller, Robert J. Nagler, Richard B. Olsen, and Olivier V. Pictet. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. 1993.

[39] Thierry Ane and Helyette Geman. Order flow, transaction clock, and normality of asset returns. *The Journal of Finance*, 55(5):2259–2284, 2000.

[40] Li Li and Zuliu Hu. Responses of the stock market to macroeconomic announcements across economic states. *IMF Working Paper*, 1998.

[41] Jean-Philippe Bouchaud. Price impact. *Encyclopedia of quantitative finance*, 2010.

[42] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity market simulation for ai research. *arXiv preprint arXiv:1904.12066*, 2019.

[43] David Byrd. Explaining agent-based financial market simulation. *arXiv preprint arXiv:1909.11650*, 2019.

[44] Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of political economy*, 101(1):119–137, 1993.

[45] Elaine Wah, Mason Wright, and Michael P Wellman. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 59:613–650, 2017.

[46] Steven Gjerstad. The competitive market paradox. *Journal of Economic Dynamics and Control*, 31(5):1753–1780, 2007.

[47] Steven Gjerstad and John Dickhaut. Price formation in double auctions. *Games and economic behavior*, 22(1):1–29, 1998.

[48] Frank McGroarty, Ash Booth, Enrico Gerding, and Venkata L. Raju Chinthalapati. High frequency trading strategies, market fragility and price spikes: an agent based model perspective. *Annals of Operations Research*, pages 1–28, 2018.

[49] Tanmoy Chakraborty and Michael Kearns. Market making and mean reversion. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 307–314. ACM, 2011.