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], Tucker 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

Market simulation is an increasingly important method for evaluating and training trading strategies and testing "what if" scenarios. The extent to which results from these simulations can be trusted depends on how realistic the environment is for the strategies being tested. As a step towards providing benchmarks for realistic simulated markets, we enumerate measurable stylized facts of limit order book (LOB) markets across multiple asset classes from the literature. We apply these metrics to data from real markets and compare the results to data originating from simulated markets. We illustrate their use in five different simulated market configurations: The first (market replay) is frequently used in practice to evaluate trading strategies; the other four are interactive agent based simulation (IABS) configurations which combine zero intelligence agents, and agents with limited strategic behavior. These simulated agents rely on an internal "oracle" that provides a fundamental value for the asset. In traditional IABS methods the fundamental originates from a mean reverting random walk. We show that markets exhibit more realistic behavior when the fundamental arises from historical market data. We further experimentally illustrate the effectiveness of IABS techniques as opposed to market replay.

KEYWORDS

Multi-agent simulations; limit order books; market microstructure

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. As a motivating example, a pension fund may have concluded that it should reduce its holdings in a particular stock by a large amount and therefore trigger a sell order for that asset. If this order was sent to an exchange as a market sell order, the price would likely fall significantly and provide the seller a less than desirable average price. In order to reduce transaction costs, it is a common practice to design *execution strategies* distributing a larger order as a set of smaller orders over time, thereby minimizing the price impact [2].

Significant research effort is aimed at applying Reinforcement Learning (RL) to such trading problems in which the learners are trained in simulation: an RL market-maker was presented in [45]; an RL approach to algorithmic execution was introduced in [40]; deep hedging a portfolio of derivatives in the presence of market friction was considered in [13]; LSTM representations for an RL trading agent are given in [38].

In these financial trading problems, the statistical properties of the environment are often unknown and difficult to model. In such cases, repeating the process a number of times in a simulated environment obviates the need to know transition probabilities and an optimal policy can be learned from the gained simulated experience, necessitating realistic market simulation tools.

In real-time algorithmic trading, the actions of any given agent incurs a response from the other market participants. In simulation, autonomous agents can choose to place orders at any time and the market's response to them will not be reflected in historical data. Therefore, simple market replay of historical orders is not sufficient for effective back testing or strategy construction. Interactive agent-based simulation (IABS) has the potential to realistically simulate the interactions between individual market participants [39], [redacted] In such simulators, prices arise 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 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 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 limit order book (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 (FIFO), is used to pair incoming buy and sell order interest [9].

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 (bid), or should not be gone

ICAIF '20, October 15–16, 2020, New York, NY, USA

ACM ISBN 978-1-4503-7584-9/20/10…$15.00
https://doi.org/10.1145/3383455.3422561

below in the case of a sell order (ask). A limit order queues a resting order in the LOB at the corresponding side of the book. Placing a limit order at a certain price level is sometimes referred to as placing a quote. 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.

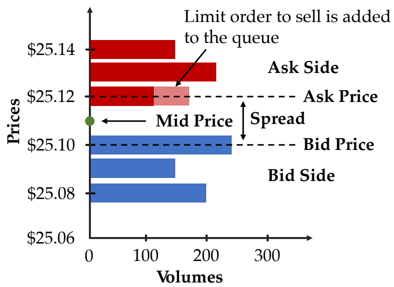

Figure 1: Visualization of the LOB structure.

1.3 Stylized facts

Many of the metrics we present below are derived from observation of the market over time. Properties of market behavior that are repeated across a wide range of instruments, markets and time periods are referred to as *stylized facts* [18].

Evaluating stylized facts for simulated data and comparing 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 market matching mechanisms has been widely discussed in the literature. If certain 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, Farmer et al. [22] showed that a market simulation that consists only of ZI agents is able to reproduce price and spread dynamics as well as market impact. The ability of ZI agents to reproduce fat tails and long range dependence was shown in [42].

In this paper, we list multiple measurable stylized facts of LOB markets known in the literature. We apply these stylized facts as a metric of realism to data from real markets and compare the results to data arising from simulated markets. We further provide five different simulated market configurations – four that are based on IABS which combine ZI agents with limited strategic behavior agents; and one market replay configuration that is commonly used in practice due to its simplicity but does not constitute a true IABS as interactions between multiple market participants are not present. Our simulated agents rely on an internal "oracle" that provides a fundamental value for the asset – we consider random and historical fundamental models. We describe the configurations and agent types in detail in Sections 3.2 and 3.3.

1.4 Related work on market simulation

Market simulation is an increasingly important technique for evaluating trading strategies and testing "what if" market scenarios. The extent to which results from such simulations can be trusted depends on how accurately they emulate real world environments.

IABS methods allow us to study phenomena that emerge as a consequence of multiple participant interactions which are difficult to model otherwise. Examples of such complex phenomena include both routine market microstructure events such as market response to an individual participant's trading [19] and rare events such as flash crashes [16, 31, 41] as well as extreme market shocks.

In the financial literature there are examples of simulators that use learning behaviors with differing views of past data [27, 32]. Wellman helped establish an empirical approach to the study of markets using simulated multi-agent systems [52] using a technique known as Empirical Game Theoretic Analysis (EGTA). Levy et al. [34] and Wah and Wellman [49] take a *synchronous* approach to simulation, wherein time is discretized between the start and end of simulation and each step is individually simulated. Jacobs et al [28, 29] introduced an *asynchronous* financial simulation framework called JLMSim. NASDAQ researchers experimentally demonstrated using IABS that under some agent scenarios reducing tick size would lead to increased spreads (an undesirable property) and would negatively impact price discovery [7, 21].

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 [2, 3]. Presence of market impact in real time implies that a realistic trading strategy simulation should include deviation from historical data. In the 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 [46]. A simple two-agent simulated market environment that consists of an algorithmic trading agent and the rest of the market with partial deviation from historical prices is presented in [48]. This model is however only suited for small order placement, and is unable to capture more complex dynamics of transient price impact [10, 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-based modeling typically relies on common sense hand-crafted rules (e.g., [42]), 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 [47] and using Bayesian techniques—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., [48, 53]). Multi-agent LOB environments can be viewed as a non-cooperative games in which every agent pursues their own goal and there is no communication between the agents [26]. Agents that learn to maximize their long term rewards by reinforcement from empirical equilibrium environments have been discussed in [44].

Other approaches to IABS realism can include inverse learning agents' rewards from the market [53]; generating synthetic LOB data using Generative Adversarial Networks [35]; incorporating feedback from real-time trading into the simulation [43] and building adaptive agents that are governed by the evolutionary principles and can learn from experience [33, 37].

2 REALISM METRICS

One way to establish IABS realism is to ensure that simulated LOB time series mimic stylized facts derived from real market histories. Below we review several groups of such stylized facts for IABS configurations, subsets of these stylized facts have been used to justify environment realism in [22, 42, 51, 53].

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 the mid-price as $m_t = \frac{a_t+b_t}{2}$. 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 the size of a new order placed into the LOB and let T be the lifetime of an order until it is fully executed or cancelled. 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 respectively. Partition the LOB price and volume time series into 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 $P(\cdot)$ denote the 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 [18] for equity markets as well as in [6] for foreign exchange and rates markets. We enumerate these statistical properties below and present them as assertions regarding the relevant data:

- **Absence of autocorrelations** Linear autocorrelations $\mathrm{corr}(r_{t+\tau,\Delta t}, r_{t,\Delta t})$ 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 slimmer 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 a 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 $\mathrm{corr}(r^2_{t+\tau,\Delta t}, r^2_{t,\Delta t})$. Empirical results on various equities indicate that this quantity remains positive over several days, which indicate periods of high volatility clustering [18].
- **Long range dependence** If one looks at the autocorrelation function of absolute returns as a function of time lag $f(\tau) = \mathrm{corr}(|r_{t+\tau,\Delta t}|, |r_{t,\Delta t}|)$, it is empirically shown that $f(\tau)$ decays according to the power law distribution $f(\tau) \sim \tau^{-\beta}$ with exponent $\beta \in [0.2, 0.4]$ [18].
- **Gain/loss asymmetry** Gain/loss asymmetry is prevalent for equity price returns as stocks lose value faster than they grow [18]. 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.
- **Volume/volatility positive correlation** Volume and volatility are positively correlated. A linear regression relationship $\mu_{V_\tau} \sim \alpha + \beta\sigma_{\tau,\Delta t}$ can be derived from the data [12].
- **Asset 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

- **Quote volumes** Aggregate quote volumes at best bid V_b (and respectively volumes at best ask V_a) are distributed according to Gamma distribution [11] for $\gamma \leq 1$, whence $P(V_b) \sim \exp^{-V_b} V_b^{-1+\gamma}$.
- **Quote sizes** Quote sizes are roughly power-law distributed [9]. For instance, Abergel et al. [1] 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 $1 + \mu \approx 2.3 - 2.7$. Deviating from power-law behaviour, orders tend to have a round number of shares (i.e. multiples of 10, 100, etc. are more common than neighboring sizes).
- **Number of quotes in a fixed time window** The number of new quotes in a fixed time window can be approximated by gamma or lognormal distributions [1].
- **Quote inter-arrival times** In the literature, LOB quote inter-arrival times are suggested to be fit into exponential [35], lognormal, and Weibull distributions [1].
- **New quote prices** Prices at which new quotes are placed, are power-law distributed around the mid price [1]. Specifically, $P(\Delta) \sim \Delta^{-(1+\mu)}$ with $1 + \mu \approx 1.6$ [11].
- **Cancellation time** 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 [1].
- **Time correlation of order flow** Individual agent's order placement decisions depend on other agents' actions [42].

2.4 Stylized facts about non-stationary patterns

- **Intraday volume patterns** LOB quote and transaction volumes are known to exhibit strong intraday patterns. For instance, historical foreign exchange trading volumes can be approximated by the "U-shaped" polynomial regional sessions that correspond to New York, London, and Tokyo trading [20]. Similarly, in most equity markets, volumes are highest at the beginning of the trading day, followed by a period of lower activity, and then spike again at the end of the trading day, which also suggests a "U-shaped" polynomial approximation [9]. Note that making a transformation from physical time to tick (or transaction) time may help adjusting for intraday non-stationarity [4] for the purpose of volume profile modeling.

- **Seasonal volume patterns** Some assets, especially those for which consumer demand 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 [36]. 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

$$P_\tau = \frac{\left|V_{\text{buy},\tau} - V_{\text{ask},\tau}\right|}{V_{\text{buy},\tau} + V_{\text{ask},\tau}}.$$

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 a relationship of the form $M_i \sim \alpha P_i^\beta$ through the data [3, 8, 22].

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 [18]. For futures, as an example, asset price moves across term structure are highly correlated and exhibits consistent patterns uncovered by Principal Components Analysis (e.g., [6]). It is also worth noting that extreme returns (e.g., 99$^{\text{th}}$-percentile returns that occur during financial crises) across various stocks or asset classes can be extremely correlated while their average returns are not [18].

3 EXPERIMENTS WITH A MULTI-AGENT SIMULATION

3.1 Simulation environment

In order to evaluate the ability of a given agent configuration to reproduce stylized facts about the market, we employ an agent-based interactive discrete event simulation environment [15]. The environment 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.

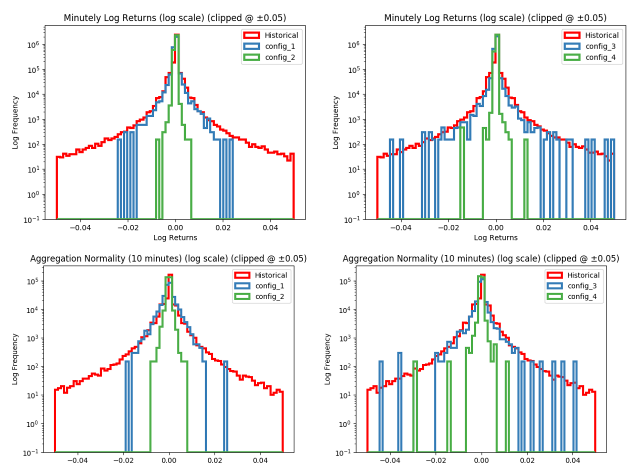

Figure 2: One-minute (top) and ten-minute (bottom) log return distributions for configs 1 and 2 (left) and configs 3 and 4 (right).

3.2 Background agents

In order to conduct experiments, we specify two ZI (value and noise agents) and two strategic agent types (market maker and momentum agents). Originally introduced in [24], the ZI agent class includes a variety of agents that do not base their trading decisions on the knowledge of LOB microstructure. They can, however, have access to the exogenous price of an asset (e.g. value agents in the implementation described below) which represents the agent's understanding of the outside world (e.g. earnings reports, macroeconomic events, etc) [50, 51]. This reference stream is called the *fundamental*. We consider two ways of modeling the fundamental price of an asset: (1) by a stochastic mean reverting process with megashock events [14] which we further call a **random** fundamental; and (2) using a historical price series – which we call a **historic** fundamental. In contrast, strategic agents' decisions are governed exclusively by LOB microstructure considerations.

Noise Agents: Noise agents are designed to simulate the action of retail traders to trade on demand (e.g. [30]). Each noise agent trades once a day by placing a market order. The direction and the size of the trade are chosen randomly. In order to model higher trading activity in the beginning and at the end of a trading day (also known as intraday volume smile), noise agent arrival time is sampled from a U-quadratic distribution over the interval $[t_{\text{open}}, t_{\text{close}}]$, where t_{open} and t_{close} are the start and end of the trading day, respectively. This arrival distribution pattern represents human trader propensity to be more active toward the beginning and the end of the trading day, as observed by [20].

Value Agents: The value agents are designed to simulate the actions of fundamental traders that trade according to their belief of the exogenous value of a stock, but without any view of the LOB microstructure. The external value of stock price is modeled by a fundamental price stream. Each value agent arrives to the market multiple times according to a Poisson process and chooses to buy or sell a stock depending on whether it is cheap or expensive relative to its noisy observation of a fundamental price. Once the side of an order is determined, the value agent places a limit order at a random level either inside the spread or deeper into the LOB. Value agents assist LOB price formation by bringing external information to the LOB and are conceptually related to informed traders widely discussed in the literature (e.g. [30]).

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. At time t, the agent starts by cancelling any of its unexecuted orders. It then looks at the LOB half-spread given by $s_t = \frac{a_t - b_t}{2}$ and mid-price $\frac{a_t + b_2}{2}$, and places new price quotes of constant size K around the mid-price at levels $m_t - s_t - N, \ldots, m_t - s_t$ and $m_t + s_t, \ldots, m_t + s_t + N$, N levels deep into the LOB. This market maker model is similar in spirit to that of [17, 50] with the distinction that out market maker does not use any reference price series to determine its mid, and rather adapts to the LOB mid price that is dictated by other market participants.

Momentum agents: The momentum agents base their trading decision on observed LOB price trends. Our implementation compares T_1 past mid-price observations to T_2 past observations with $T_1 < T_2$ and places a buy order of random size, if the former exceeds the latter and a sell order otherwise. Finally, the momentum agents are configured to arrive to the market at a constant rate.

Market replay agent: The market replay agent works by passing historical orders for a given day and security into the simulation exchange's matching engine. At any given time in the simulation, absent any other trading agents, the order book exactly replicates the historical order book at that time. Figure 3 illustrates the main limitation of market replay: as opposed to the IABS scenario wherein an experimental agent interacts with multiple diverse market participants via an exchange, the market replay scenario only allows replay of historical data — the market does not respond to the actions of the experimental agent, although the experimental agent responds to the actions of the market.

3.3 Agent configurations

We present four different agent configurations as described in Table 1 and further evaluate them for realism. Our reasoning for

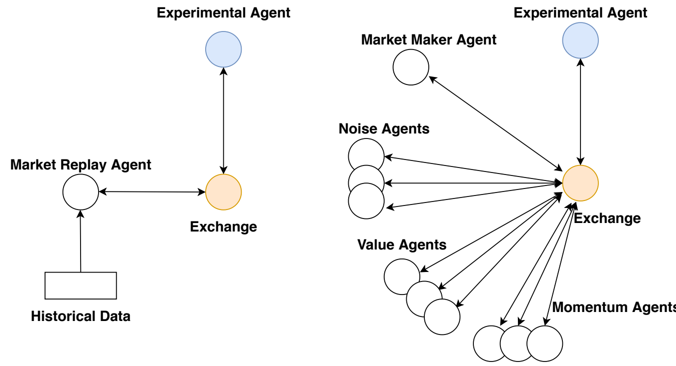

Figure 3: Market replay (left) vs. IABS (right).

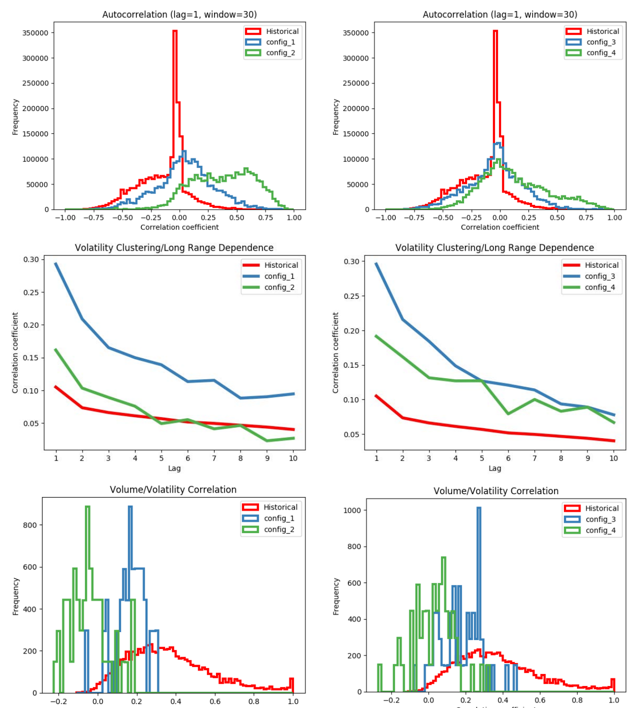

Figure 4: Stylized facts for configs 1 and 2 (*left*), and configs 3 and 4 (*right*). (*top*) Distributions of return autocorrelation. (*center*) Average autocorrelation of square returns as a function of time lag. (*bottom*) Volume/volatility correlation distributions.

introducing progressively more complex configuration design is the following. Since it is known that some stylized facts can be modeled exclusively by agents in the ZI family (i.e. [22]) we want to introduce minimal strategic behavior on top of zero intelligence, in order to gain insights into which LOB properties we can reproduce and which will require more sophisticated agent modeling (possibly by introducing learning agents).

An additional dimension that we want to assess are random and historic approaches to modeling fundamental time series for the purpose of simulated realism. We expect that using a historic fundamental will be more realistic when looking at price returns, and want to evaluate the degree of agent diversity that is needed for the use of historic price fundamental to be realistic in terms of other stylized facts.

Finally, the market replay agent configuration is expected to be most realistic in terms of stylized facts by definition. However, market replay does not constitute a true IABS as interactions between multiple agents are not present, as depicted in Figure 3.

Config	Agents	Fundamental	ZI/Strategic
1	5000 noise, 100 value	Random	ZI
2	5000 noise, 100 value, 1 market maker, 25 momentum	Random	Both
3	5000 noise, 100 value	Historic	ZI
4	5000 noise, 100 value, 1 market maker, 25 momentum	Historic	Both
5	1 market replay	Historic	not IABS

Table 1: Agent configurations.

4 EXPERIMENTAL RESULTS

To assess the realism of the first four IABS configurations, we examine stylized facts relating to asset returns and order flow described in Sections 2.2 and 2.3 respectively.

4.1 Stylized facts about asset return distributions

To derive historical asset return distributions, we analyze minutely intraday log returns of 30 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.**
  Config 3 is the most realistic as it exhibits the heaviest tails. This realism can be attributed to config 3 being closest to historical data. The addition of a market maker appears to reduce volatility of returns, as seen with configs 2 and 4, making them less realistic (Figure 2).
- **Absence of autocorrelations** Config 3 is the most realistic. Config 4 is also realistic but to a lesser degree. (Figure 4).
- **Volatility clustering.** The average autocorrelation of square returns decays for both historical and simulated data in all configurations as time lag increases. The volatility clustering of config 2 is most similar to historical data (Figure 4).
- **Volume/volatility correlation.** None of the configs exhibit volume/volatility correlation properties similar to historical data. (Figure 4).

4.2 Stylized facts about volumes and order flow

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

- **Number of quotes in a fixed time window.** Figure 5 (*top*) show the distribution for number of quotes occuring in a five-minute window for the simulated vs. the historical data. The number of quotes is normalized to the total daily quotes for easier visual comparison. We find that gamma distribution produces a good fit for all curves.
- **Intraday quote volume patterns.** Quadratic curves have been fitted to the quote volume in a given 5-,minute period throughout the day to demonstrate the "U-shaped" pattern of historical intraday quote volumes (see Figure 5 (*center*)). The data was normalized to the mean volume. We find the simulated data reproduces the volume “smile” well.
- **Quote interarrival times.** Figure 5 (*bottom*) show the distribution of quote interarrival times for simulated data as well as the historical data. The historical data fits a Weibull distribution well. Large support at zero interarrival time leads to poor fit for the simulated data in all configurations.

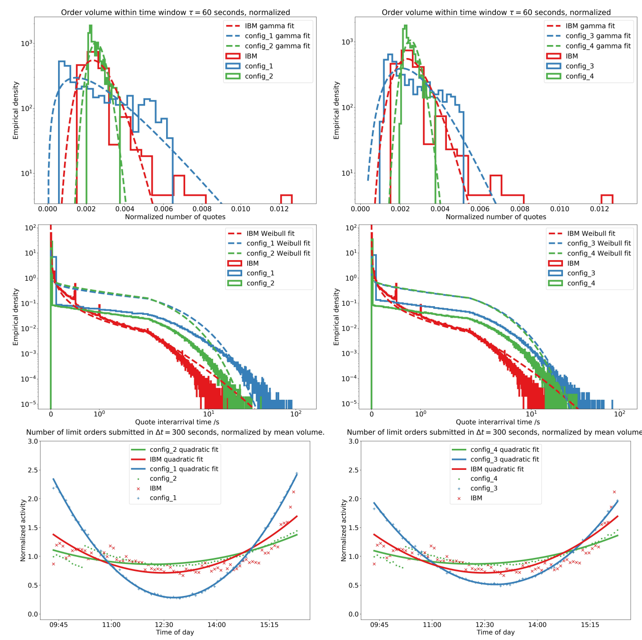

Figure 5: Order flow stylized facts for configs 1 and 2 (*left*) and configs 3 and 4 (*right*) as compared with empirical data (*top*) Number of quotes in a five-minute window. (*center*) Quote interarrival times. (*bottom*) Intraday quote volume profiles.

4.3 Stylized facts realism: summary

In summary, configurations 2 and 4 with diverse mix of ZI and strategic agent populations can replicate stylized facts about order flows better. This agrees with our intuition as in reality many groups of market agents react to limit order book information.

Configuration 3 can replicate stylized facts about asset return distributions best as value agents in this configuration are driven by historic fundamental and it is closest to historic market replay. We, however, want to underline that the agent homogeneity in configuration 3 is likely to be a limitation for realistic order book simulations as our analysis of order flow stylized facts illustrates.

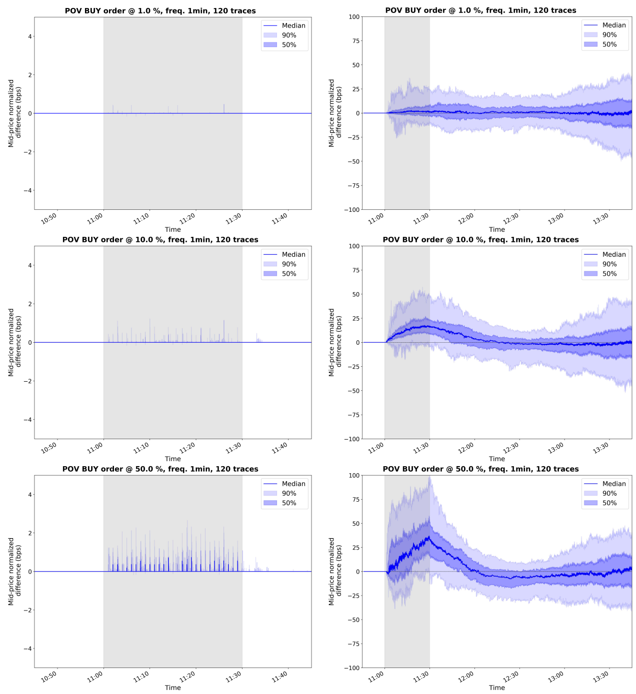

Figure 6: Market replay (*left*) and Configuration 4 (*right*) - with and without POV execution agent, executing at 1% (*top row*), 10% (*middle row*) and 50% (*bottom row*) of transacted volume at 1 minute intervals between 11:00 AM and 11:30 AM. Market replay x-axis is truncated and y-axis expanded for clarity. The thick blue line represents the median, with 50% and 90% quantiles indicated by the shaded regions.

4.4 Market impact experiment

We design an execution experiment that demonstrates that despite the use of historical fundamental by the value agents, IABS in configuration 4 is able to simulate realistic market impact dynamics of a large order execution. In this experiment, we compare market replay to interactive agent-based simulation (IABS) in configuration 4 for evaluating the price impact of limit order placement in a simulated market.

In particular, we carry out an experiment simulating a percentage of volume (POV) strategy typically used in practice. A POV algorithm is defined by a percentage level $\zeta \in (0, 1]$, a wake-up period, a direction (buy or sell) and a target quantity. At each wake-up, the total transacted volume since the previous wake-up V_τ is computed, then a market order of size $\zeta \cdot V_\tau$ is placed in the specified direction. This continues until the trader has reached the target quantity. In our experiment, we place POV orders for $\zeta \in \{0.01, 0.1, 0.5\}$ between 11:00AM and 11:30AM for both market replay and configuration 4[1], using 120 random seeds and 20 dates for IBM for config 3 with the same dates in market replay.

Figure 6 visualizes the price impact of order execution that starts at 11:00AM with the arrival of the POV order and continues for the duration of an order until 11:30 AM. In both market replay and configuration 4 (IABS), the mid-price deviates from the baseline configuration without the POV order. The key difference is in how the mid-price evolves after the first order is placed. In the case of market replay, spikes in the normalized difference between the mid-prices are observed as a result of the new orders in the historical order book being replayed. IABS, on the other hand, allows for the adaptation of the agents to the new observed mid-price which results in a different mid-price evolution that is not reproduced using market replay only. Note that in the IABS configuration, the presence of price impact remains past the end of the order execution at 11:30 AM which aligns with permanence of market impact described in [2, 5], demonstrating behavior that is not reproducible using market replay only. We note also that the price impact increases as a function of the percentage level ζ, as is observed in real markets.

[1]For this experiment the arrival rate of the value agents was reduced by a factor ≈ 1.4 as a more realistic impact trace was desired.

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 evaluated four experimental configurations of agent types—ZI agents that follow random fundamental, ZI agents that follow historic fundamental, ZI and strategic agents that follow a random fundamental and ZI and strategic agents that follow a historic fundamental. We found that a diverse mix of ZI and minimal strategic behavior agent types and use of historic fundamental result in more realistic LOB simulation. We also demonstrated that using historic time series as a fundamental reference stream for the value agents is not a limiting factor in simulating emergent phenomena (such as market impact of trading) and allows simulated price series to diverge significantly from their historic trajectories.

Although we observed that the configurations with more diverse ZI and strategic agent populations and historic fundamental lead to statistics that more closely mimic real markets, we acknowledge that there is much room for improvement. In particular, order flow correlation and interarrival time properties were not well reproduced. Related literature suggests that 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 the introduction of online learning agents [42]. For example, ref. [33] details comparisons of non-learning and learning agents and concludes that agents capable of learning and adaption to other agent flows are better able to replicate stylized facts about long range dependence and correlation between volume and volatility. Moreover, in real markets, rational agents evolve over time by learning to expand effective and cull ineffective trading strategies [37]. Hence, we believe that furnishing autonomous LOB agents with the ability to learn from experience will be a step towards making simulated environments more robust.

6 ACKNOWLEDGEMENTS

This paper was prepared for informational purposes ["in part" if the work is collaborative with external partners] by the Artificial Intelligence Research group of JPMorgan Chase and 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 representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information 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, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

REFERENCES

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

[2] Robert Almgren and Neil Chriss. 1999. Optimal execution of portfolio transactions. *Journal of Risk* (1999).

[3] Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. 2005. Direct Estimation of Equity Market Impact. *RISK* 18 (04 2005).

[4] Thierry Ane and Helyette Geman. 2000. Order Flow, Transaction Clock, and Normality of Asset Returns. *The Journal of Finance* 55, 5 (2000), 2259–2284.

[5] Tucker Hybinette Balch, Mahmoud Mahfouz, Joshua Lockhart, Maria Hybinette, and David Byrd. 2019. How to Evaluate Trading Strategies: Single Agent Market Replay or Multiple Agent Interactive Simulation? *arXiv preprint arXiv:1906.12010* (2019).

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

[7] Eric Bonabeau. 2002. 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).

[8] Jean-Philippe Bouchaud. 2010. Price impact. *Encyclopedia of quantitative finance* (2010).

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

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

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

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

[13] H Buehler, L Gonon, J Teichmann, and B Wood. 2018. Deep Hedging. *arXiv:1802.03042* (2018).

[14] David Byrd. 2019. Explaining Agent-Based Financial Market Simulation. *arXiv preprint arXiv:1909.11650* (2019).

[15] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. 2019. ABIDES: Towards High-Fidelity Market Simulation for AI Research. *arXiv preprint arXiv:1904.12066* (2019).

[16] CFTC. 2010. 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).

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

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

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

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

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

[22] J Farmer, Paolo Patelli, and Ilija Zovko. 2005. The Predictive Power of Zero Intelligence in Financial Markets. *Proceedings of the National Academy of Sciences of the United States of America* 102 (03 2005), 2254–9.

[23] Jim Gatheral. 2010. No-Dynamic-Arbitrage and Market Impact. *Quantitative Finance* (2010).

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

[25] Jakob Grazzini, Matteo Richiardi, and Efthymios Tsionas. 2015. Bayesian Estimation of Agent-Based Models. *Journal of Economic Dynamics and Control* (11 2015).

[26] Junling Hu and Michael P. Wellman. 1998. Multiagent Reinforcement Learning: Theoretical Framework and an Algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML '98)*.

[27] Bruce I. Jacobs, Kenneth N. Levy, and Harry M. Markowitz. 2004. Financial Market Simulation. *The Journal of Portfolio Management* 30, 5 (2004), 142–152. https://doi.org/10.3905/jpm.2004.442640 arXiv:https://jpm.pm-research.com/content/30/5/142.full.pdf

[28] Bruce I Jacobs, Kenneth N Levy, and Harry M Markowitz. 2004. Financial market simulation. *The Journal of Portfolio Management* 30, 5 (2004), 142–152.

[29] Bruce I Jacobs, Kenneth N Levy, and Harry M Markowitz. 2010. Simulating security markets in dynamic and equilibrium modes. *Financial Analysts Journal* 66, 5 (2010), 42–53.

[30] Albert S Kyle. 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society* (1985), 1315–1335.

[31] Sandrine Leal and Mauro Napoletano. 2016. 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).

[32] B. LeBaron. 2001. A builder's guide to agent-based financial markets. *Quantitative Finance* 1, 2 (2001), 254–261.

[33] Blake LeBaron and Ryuichi Yamamoto. 2007. Long-Memory in an Order-Driven Market.

[34] Moshe Levy, Haim Levy, and Sorin Solomon. 1994. A microscopic model of the stock market: cycles, booms, and crashes. *Economics Letters* 45, 1 (1994), 103–111.

[35] Junyi Li, Xintong Wang, Yaoyang Lin, Arunesh Sinha, and Michael P. Wellman. 2019. Generating Realistic Stock Market Order Streams. https://openreview.net/forum?id=rke41hC5Km

[36] Li Li and Zuliu Hu. 1998. Responses of the Stock Market to Macroeconomic Announcements Across Economic States. *IMF Working Paper* (1998).

[37] Andrew Lo. 2004. Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting* 7 (01 2004).

[38] David Lu. 2017. Agent Inspired Trading Using Recurrent Reinforcement Learning and LSTM Neural Networks. *arXiv:1707.07338* (07 2017).

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

[40] Y. Nevmyvaka, Y. Feng, and M. Kearns. 2006. Reinforcement Learning for Optimized Trade Execution. In *Proceedings of the 23rd International Conference on Machine Learning*. Pittsburg, PA.

[41] Mark Paddrik, Roy Hayes, Andrew Todd, Steve Yang, Peter Beling, and William Scherer. 2012. 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)*. IEEE, 1–8.

[42] Imon Palit, Steve Phelps, and Wing Lon Ng. 2012. Can a Zero-Intelligence Plus Model Explain the Stylized Facts of Financial Time Series Data? 653–660.

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

[44] L. Julian Schvartzman and Michael P. Wellman. 2009. 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*.

[45] T. Spooner, J. Fearnley, R. Savani, and A. Koukorinis. 2018. Market Making via Reinforcement Learning. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*. Stockholm, Sweden, 434–442.

[46] 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*.

[47] Gabriele Tedeschi, Maria Recchioni, and Mauro Gallegati. 2013. A Calibration Procedure for Analyzing Stock Price Dynamics in an Agent-Based Framework. *http://ssrn.com/abstract=2330739* (09 2013).

[48] Svitlana Vyetrenko and Shaojie Xu. 2019. Risk-Sensitive Compact Decision Trees for Autonomous Execution in Presence of Simulated Market Response. In *ICML 2019 Workshop on AI in Finance*.

[49] 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*. ACM, 855–872.

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

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

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

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