

Deep Learning Meets Queue-Reactive: A Framework for Realistic Limit Order Book Simulation

Hamza Bodor^{a,b}, Laurent Carlier^b

^a*Université Paris 1 Panthéon-Sorbonne, Centre d'Economie de la Sorbonne, 106 Boulevard de l'Hôpital, Paris Cedex 13, 75642, France*

^b*BNP Paribas Corporate and Institutional Banking, Global Markets Data & Artificial Intelligence Lab, 20 boulevard des Italiens, Paris, 75009, France*

Abstract

The Queue-Reactive model introduced by [Huang et al. \(2015\)](#) has become a standard tool for limit order book modeling, widely adopted by both researchers and practitioners for its simplicity and effectiveness. We present the Multidimensional Deep Queue-Reactive (MDQR) model, which extends this framework in three ways: it relaxes the assumption of queue independence, enriches the state space with market features, and models the distribution of order sizes. Through a neural network architecture, the model learns complex dependencies between different price levels and adapts to varying market conditions, while preserving the interpretable point-process foundation of the original framework. Using data from the Bund futures market, we show that MDQR captures key market properties including the square-root law of market impact, cross-queue correlations, and realistic order size patterns. The model demonstrates particular strength in reproducing both conditional and stationary distributions of order sizes, as well as various stylized facts of market microstructure. The model achieves this while maintaining the computational efficiency needed for practical applications such as strategy development through reinforcement learning or realistic backtesting.

Keywords: Limit order book, Generative models, Market simulators, Stylized facts, Market microstructure

Email addresses: bodor.hamza@gmail.com (Hamza Bodor), laurent.carlier@bnpparibas.com (Laurent Carlier)

1. Introduction and Motivation

Financial markets have evolved significantly over the past decades, with electronic limit order books (LOBs) becoming the predominant mechanism for price formation and trading across major exchanges worldwide. [Gould et al. \(2013\)](#) provide a comprehensive overview of these markets, where participants interact by submitting limit orders (indicating willingness to trade at specific prices) and market orders (executing against existing limit orders), creating a complex and high-dimensional system whose dynamics emerge from the interactions of diverse market participants ([Cont, 2011](#)).

Understanding and modeling these dynamics is crucial for several reasons. First, market participants need realistic simulators to develop and test trading strategies before deploying them in live markets, where testing can be both risky and costly ([Cartea et al., 2015](#)). Second, regulators require tools to assess market stability and evaluate potential policy changes, as demonstrated by [Kirilenko et al. \(2017\)](#) in their analysis of the Flash Crash. Third, researchers seek to understand the fundamental mechanisms driving price formation and market impact ([Gatheral, 2010](#)).

The complexity of limit order markets poses significant modeling challenges. Order arrivals exhibit complex temporal dependencies, with order flow at one price level potentially influencing dynamics at other levels. Order sizes vary considerably and show distinctive patterns. Furthermore, the market’s response to large trades manifests through intricate feedback mechanisms, resulting in well-documented phenomena such as the “square-root law” of market impact ([Gatheral, 2010](#); [Cont and De Larrard, 2014](#)).

Traditional approaches to modeling LOBs often rely on simplifying assumptions to maintain tractability. [Huang et al. \(2015\)](#) made a significant advance with their Queue-Reactive (QR) model by capturing the dependency between order flow and queue sizes while maintaining analytical tractability. Recently, [Bodor and Carlier \(2024a\)](#) extensively investigated the performance of the QR model against multiple stylized facts of the Bund Market and proposed several extensions, particularly focusing on order size modeling. However, these enhanced models still maintain strong independence assumptions between price levels and simplified order size distributions.

The emergence of deep learning techniques offers new possibilities for capturing these complex dependencies while maintaining computational feasibility. Recent successes in applying neural networks to financial markets, as demonstrated by [Sirignano and Cont \(2019\)](#), suggest that these methods could help overcome the limitations of traditional approaches while preserving their interpretability.

In this paper, we propose a deep learning extension of the Queue-Reactive model that significantly increases its flexibility while maintaining its economic interpretability. Our approach combines the theoretical insights of the QR framework with the representational power of deep neural networks. The key contributions of our work include:

1. A neural network extension of the Queue-Reactive framework for order book modeling balancing interpretability and computational efficiency;
2. Integration of cross-price level dependencies and realistic order size modeling;
3. Empirical validation demonstrating the model’s ability to reproduce sophisticated market impact patterns and key stylized facts;

4. A framework that bridges the gap between traditional stochastic models and modern machine learning approaches.

The remainder of this paper is organized as follows. Section 2 reviews the literature on limit order book modeling. Section 3 presents the Queue-Reactive model, its extensions, and limitations. Section 4 introduces the MDQR framework and analyzes its properties. Section 5 compares our approach with related work. Section 6 concludes with directions for future research.

2. Related Work

The modeling of limit order book dynamics has attracted significant attention from both academics and practitioners, leading to diverse approaches ranging from purely statistical models to sophisticated machine learning techniques. Early empirical studies by [Bouchaud et al. \(2002\)](#) and [Potters and Bouchaud \(2003\)](#) established key stylized facts about order flow, price formation, and market impact that subsequent models sought to reproduce.

Statistical approaches, particularly those based on point processes, formed the initial framework for modeling order book dynamics. [Smith et al. \(2003\)](#) and [Cont et al. \(2010\)](#) introduced simple Poisson process models that provided valuable analytical insights while maintaining tractability, though they failed to capture the complex dependencies observed in empirical data. This limitation led to the development of more sophisticated models using Hawkes processes, as demonstrated by [Abergel and Jedidi \(2013, 2015\)](#), which better captured the clustering of events and temporal dependencies in order flow. A significant advancement came when [Huang et al. \(2015\)](#) introduced the Queue-Reactive model, making arrival intensities dependent on queue sizes and better reflecting the empirical relationship between order flow and market state. Recently, [Bodor and Carlier \(2024a\)](#) extended this framework by incorporating additional market features, demonstrating improved empirical fit.

[Paddrik et al. \(2012\)](#) offered a different perspective through agent-based models focusing on the behavior of market participants. While these models provided valuable insights into market mechanisms, particularly in stress scenarios like the Flash Crash studied by [Kirilenko et al. \(2017\)](#), the challenge of calibrating multiple agent behaviors to match empirical data has limited their practical application ([Axtell and Farmer, 2022](#); [Vyetrenko et al., 2020](#)).

The emergence of machine learning techniques has opened new avenues for LOB modeling. [Sirignano and Cont \(2019\)](#) demonstrated the power of deep learning in forecasting price movements, while recent work has focused on generative modeling of entire order book dynamics. Notably, [Coletta et al. \(2022\)](#) developed a GAN-based approach for LOB simulation, focusing on explainability and robustness. Their work provided frameworks for evaluating generator quality and ensuring stable training, critical considerations for practical applications. Building on this foundation, [Cont et al. \(2023\)](#) extended the GAN-based approach, successfully reproducing complex market behaviors, including the square-root law of market impact. A particularly significant contribution came from

Hultin et al. (2023), who adopted a recurrent neural network-based approach that successfully captured temporal dependencies while preserving key statistical properties of order book dynamics.

Recent work has also explored hybrid approaches that combine traditional modeling frameworks with machine learning techniques. Cont and Müller (2021) modeled the LOB as a density evolving according to a stochastic partial differential equation, while using neural networks to capture complex dependencies in the driving processes. These hybrid approaches suggest the potential for combining the interpretability of traditional models with the flexibility of modern machine learning methods.

Our work builds on these developments by proposing a deep learning extension of the Queue-Reactive model that maintains its economic interpretability while significantly increasing its flexibility to capture complex market dynamics. This approach bridges the gap between traditional stochastic models and modern machine learning techniques, offering a framework that benefits from both traditions.

3. The Queue-Reactive Model and Its Extensions

The Queue-Reactive (QR) model, introduced by Huang et al. (2015), provides an elegant framework for modeling limit order book dynamics. At its core, the model views the limit order book as a collection of queues centered around a reference price, typically the mid-price. Each queue represents the accumulated volume of orders at a specific price level, with the book divided into bid and ask sides. The evolution of these queues is driven by three types of events: limit orders that add to existing queues, cancellations that remove orders, and market orders that execute against available liquidity.

What distinguishes the QR model from earlier approaches is its central insight: the arrival rates of different order types depend on the current state of the queues. This state dependency captures a key empirical observation - market participants adjust their behavior based on visible order book conditions. For instance, a large queue at the best bid price might attract more market sell orders, while simultaneously discouraging additional limit buy orders at that level.

The model represents the limit order book as a $2K$ -dimensional vector, where K denotes the number of price levels tracked on each side of the book. The reference price p_{ref} serves as the center point, dividing the book into bid side queues $[Q_{-i} : i = 1, \dots, K]$ and ask side queues $[Q_i : i = 1, \dots, K]$. Here, $Q_{\pm i}$ represents the queue size at a distance of $i - 0.5$ ticks from the reference price, with negative indices for the bid side and positive for the ask side.

The mathematical formulation of the model captures this intuition through a system of interacting point processes. Let us consider a sequence of events occurring at a given queue $\mathcal{E} = \{e_k\}_{k=1}^N$, where each event e_k is characterized by its type $\eta_k \in \{L, C, M\}$ (limit, cancel, or market), the queue size q_k just before the event, and the time interval Δt_k since the previous event.

The key component of the model is the specification of intensity functions $\lambda^\eta(q)$ for each event type η , which determine the arrival rates of different orders based on the

current queue size. These intensities combine to form a global intensity, which is the intensity of queue update (regardless of which event is updating):

$$\Lambda(q) = \lambda^L(q) + \lambda^C(q) + \lambda^M(q).$$

The model’s likelihood function then takes the form:

$$\mathcal{L}(\{\lambda^\eta\} | \mathcal{E}) = \prod_{k=1}^K e^{-\Lambda(q_k) \Delta t_k} \lambda^{\eta_k}(q_k).$$

The maximizer of this likelihood can be derived analytically and has the form:

$$\hat{\lambda}^\eta(n) = \frac{\#\{e_k \in \mathcal{E} \mid \eta_k = \eta, q_k = n\}}{\#\{e_k \in \mathcal{E} \mid q_k = n\}} \cdot \left(\frac{1}{\#\{k \mid q_k = n\}} \sum_{\{k \mid q_k = n\}} \Delta t_k \right)^{-1} \quad (1)$$

This formulation has proven remarkably successful in reproducing several key features of limit order markets, including the relationship between queue sizes and order flow, and the mean-reversion behavior of queue sizes. However, the base model makes several simplifying assumptions that limit its ability to capture more complex market dynamics. Most notably, it assumes independence between price levels and constant order sizes within each level.

In [Bodor and Carlier \(2024a\)](#), we extend the Queue-Reactive model by incorporating order sizes into the intensity functions, modifying them to take the form:

$$\hat{\lambda}^{\eta,s}(n) = \frac{\#\{e_k \in \mathcal{E} \mid \eta_k = \eta, s_k = s, q_k = n\}}{\#\{e_k \in \mathcal{E} \mid q_k = n\}} \cdot \left(\frac{1}{\#\{k \mid q_k = n\}} \sum_{\{k \mid q_k = n\}} \Delta t_k \right)^{-1} \quad (2)$$

This extension, known as SAQR, achieves market simulations with volatility levels closer to historical data while preserving other important properties like the distribution of order sizes. However, both the original QR model and SAQR rely solely on queue sizes as features for order arrival modeling, limiting their ability to reproduce certain important stylized facts, such as excitation between events.

These limitations motivate our extension of the QR framework to accommodate a broader state space. By leveraging deep learning techniques, we can maintain the model’s interpretable structure while incorporating more realistic features such as cross-price level dependencies and varying order sizes.

3.1. Deep Queue-Reactive Model

While the QR model captures queue dynamics through its dependency on queue sizes with the likelihood function

$$\mathcal{L}(\{\lambda^\eta\} | \mathcal{E}) = \prod_{k=1}^N e^{-\Lambda(q_k) \Delta t_k} \lambda^{\eta_k}(q_k),$$

market participants’ decisions depend on broader market conditions. The Deep Queue-Reactive (DQR) model generalizes this framework by replacing the queue-size dependency q_k with a flexible state vector x_k that can incorporate any relevant market features.

The model parameterizes the intensity functions through neural networks $\lambda_\theta^\eta(x_k)$, where θ represents the network parameters. Taking the logarithm of the likelihood:

$$l(\lambda_\theta | \mathcal{E}) = \sum_{k=1}^N \log \lambda_\theta^{\eta_k}(x_k) - \Lambda_\theta(x_k) \Delta t_k$$

where $\Lambda_\theta(x_k) = \sum_{\eta \in \{L, C, M\}} \lambda_\theta^\eta(x_k)$.

The model is calibrated by minimizing the negative log-likelihood:

$$\text{loss}(\theta | \mathcal{E}) = -l(\lambda_\theta | \mathcal{E}) = \sum_{k=1}^N \Lambda_\theta(x_k) \Delta t_k - \log \lambda_\theta^{\eta_k}(x_k)$$

This neural network parameterization enables learning complex relationships between market conditions and order flow, while the flexibility in defining x_k allows incorporation of any relevant market features beyond queue sizes. Section 3.4 demonstrates how different choices of state space variables affect the simulated market properties.

3.2. Data

The calibration data is derived from Euro-Bund Futures (FGBL), specifically covering active trading days over three months from March to June 2022 for futures maturing in June 2022 (FGBLM2). The dataset is filtered to include periods from 9 AM to 6 PM, with a maximum depth of five levels on each side of the order book.¹

Level	#L ($\times 10^6$)	#C ($\times 10^6$)	#M ($\times 10^4$)	AES	AIT (ms)
1	32.9	30.1	212	6.25	57.3
2	6.62	5.63	0.51	5.04	293
3	4.59	4.58	0	6.16	392
4	2.79	2.83	0	5.69	537
5	3.93	3.63	0	3.39	443

Table 1: Descriptive statistics about events for each level.

Table 1 presents statistical details of the dataset used for calibrating the model. It includes the count of observations per event type, the Average Event Size (AES), and the Average Inter-event Time (AIT) for each queue. The statistics importantly highlight the

¹Euro-Bund Futures are derivative contracts tied to 10-year German government bonds (Bund) and traded on the Eurex Exchange. Each contract, featuring a notional coupon rate of 6%, is designed to mature between 8.5 and 10.5 years on the delivery day.

frequent occurrence of events impacting queues close to the mid-price, with a particular emphasis on the best prices.

The dataset also comprises updates to the limit order book along with the daily trade history. For model training, it is essential to focus on the order flow, which captures the sequence of events in the limit order book. We meticulously constructed the order flow by identifying whether each update resulted from liquidity provision (through a limit order) or liquidity consumption (via cancellations and market orders). More information on this process can be found in Bodor and Carlier (2024b). Moreover, a matching engine was developed to align the orders generated by each model, which were then reflected in a simulated order book.

The data undergoes further preprocessing according to the methodology outlined in Huang et al. (2015). Specifically:

- Queue sizes are adjusted based on the average event size. If the size of queue i is denoted by q_i and its average event size as AES_i , then the revised queue size is calculated as $q_i \leftarrow \lceil \frac{q_i}{AES_i} \rceil$.
- The data is segmented into periods where the reference price remains constant. If there is a change in the reference price, the interarrival time counter for events in all queues is reset to zero.
- The parameter controlling reference price movements is set to $\theta = 0.7$. This probability has been adjusted to replicate the observed ratio between price continuations and alternations.

For the calibration of model intensities and subsequent limit order book simulations, we take the following steps:

- Use the calibrated intensities to determine both the timing and the type of new orders.
- For modeling price changes, we adhere to the specified approach but avoid redrawing the limit order book from its stationary distribution to focus on endogenous market behavior.
- When the reference price changes, the size of the queue for new prices is sampled from their stationary distributions.

3.3. Neural Architecture and Training Methodology

The model employs a neural network architecture to calibrate parametric intensities across price levels, balancing predictive performance with computational requirements. The architecture consists of a Multi-Layer Perceptron (MLP) with two hidden layers of dimensions 128 and 32 neurons using hyperbolic tangent (\tanh) activations, followed by an output layer of dimension 3 with Rectified Linear Unit (ReLU) activation that produces arrival intensities for limit orders, cancellations, and market orders. This configuration

emerged from empirical testing, offering an effective balance between model expressiveness and computational efficiency—an important consideration for the simulation framework.

To enhance the model’s learning capabilities and stability, we incorporate batch normalization layers between successive dense layers. In the cases when the input vector x_k contains categorical features, we implement learnable embedding layers of dimension 2, allowing the model to discover meaningful representations of discrete variables while maintaining a compact parameter space.

The training protocol follows the subsequent methodology:

- **Data partitioning:** 80% of the dataset is allocated for training, with the remaining 20% reserved for validation and early stopping criteria
- **Optimization:** Implementation of the Adam optimizer with cyclic learning rate scheduling, oscillating between 10^{-5} and 10^{-3}
- **Early stopping:** Training termination after 10 consecutive epochs without improvement in validation loss

Table 2 summarizes the key architectural components and training parameters:

Table 2: Neural Network Architecture and Training Specifications

Component	Specification
Hidden Layers	[128, 32]
Output Dimension	3 (intensity values)
Embedding Dimension	2 (categorical features)
Optimizer	Adam with cyclic LR
Learning Rate Range	10^{-5} to 10^{-3}

This architecture demonstrates robust convergence properties while maintaining the computational efficiency necessary for practical deployment in high-frequency market scenarios.

3.4. Examples of Applications

3.4.1. Excitation between Events

In financial markets, the type of the previous event (η_{k-1}) often influences the type and intensity of the next event. This phenomenon, known as excitation, is critical for realistic modeling of event dynamics. Traditional Queue-Reactive (QR) and Size-Aware Queue-Reactive (SAQR) models fail to capture excitation because they do not explicitly model any information about the order flow history, resulting in inaccurate excitation values.

To address this, we incorporated η_{k-1} as a feature in the state space $x_k = [q_k, \eta_{k-1}]$ for the Deep Queue-Reactive (DQR) model. This inclusion enables the model to learn how the type of the last event influences subsequent events.

Figure 1 compares the transition matrices of events:

- **Historical data:** The ground truth, showing distinct excitation patterns.
- **SAQR model:** Rows are uniform, indicating the absence of excitation.
- **DQR model:** Captures realistic excitation patterns, closely matching historical data.

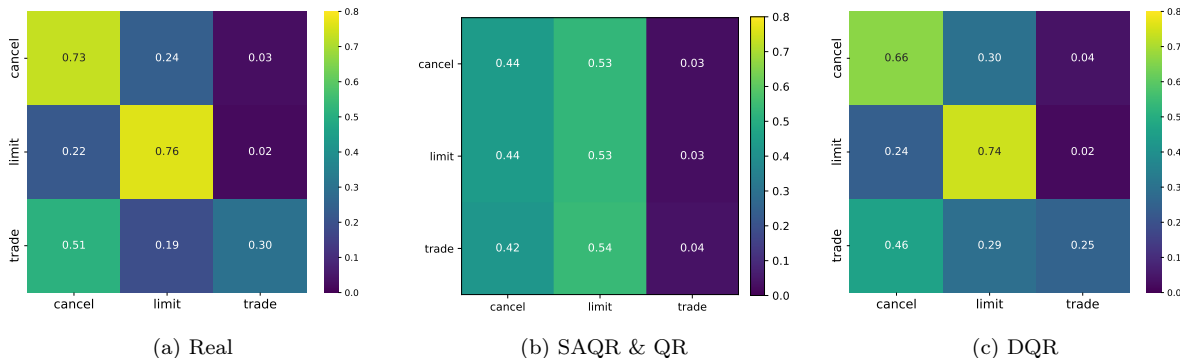


Figure 1: Transition matrix of events on simulated markets vs historical data. The rows of the matrix represent the conditional probabilities of observing a specific event type given the type of the previous event (intra-sides).

As seen in Figure 1, the DQR model reproduces excitation patterns observed in historical data, validating the importance of including η_{k-1} in the state space. In contrast, the SAQR and QR models fail to capture this key dynamic, resulting in uniform transition matrices, meaning absence of any event type interdependencies. Previous work by Wu et al. (2019) successfully reproduced these excitation patterns by developing a hybrid framework combining Hawkes processes with the QR model. The DQR framework achieves comparable results capturing these dynamics by incorporating historical event information in the state space.

3.4.2. Intraday seasonality

Another limitation of the QR model is its uniform treatment of market behavior across different times of the day. In practice, markets exhibit distinct activity regimes throughout the trading session, with heightened activity during opening and closing periods and calmer phases around lunch break, as documented in Bodor and Carlier (2024b). To capture these intraday patterns, we incorporate temporal information into the model’s state space.

To capture intra-day seasonality, we included the *hour of the last event* h_k as an additional feature in the state space x_k . This transforms x_k for each event e_k into:

$$x_k = [q_k, h_k],$$

where q_k is the size of the queue immediately before the event, and h_k is the current hour of the day (as a categorical feature).

Including h_k enables the model to learn intra-day patterns, such as variations in event intensities across trading hours (e.g., heightened activity at market open and close). Figure 2 illustrates how incorporating h_k improves the ability of the DQR model to reproduce observed intra-day seasonality in market activity.

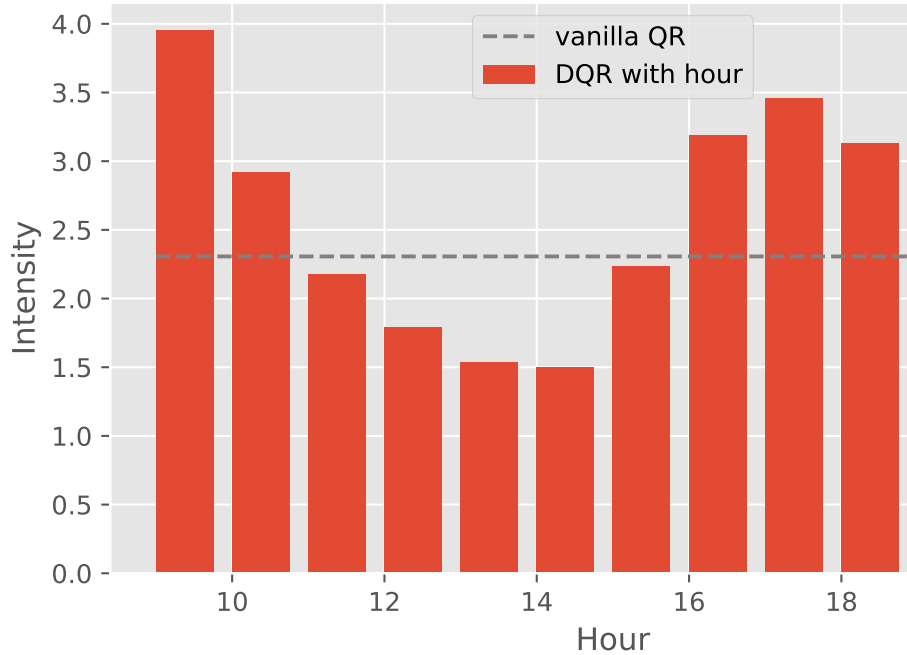


Figure 2: Comparison of DQR (with $x_k = [q_k, h_k]$) and QR trade order arrival intensities across trading hours, averaged over queue sizes.

3.4.3. Impact of Feature Enrichment on Model Performance

To further validate the impact of enriching the state space x_k , we evaluated the performance of the Deep Queue-Reactive (DQR) model across three key metrics:

- **Log-likelihood of the model:** A higher log-likelihood indicates better alignment of the model’s predictions with observed market events.
- **Balanced accuracy of next-event prediction:** The ability of the model to correctly classify the type of the next event (η_k).
- **Relative difference in time to the next event:** Measures how well the model predicts the timing of the next event compared to the actual observed values.

The results, shown in Figure 3, demonstrate the performance improvements achieved by incorporating additional features into x_k , such as the hour of the last event (h_k) and the type of the last event (η_{k-1}).

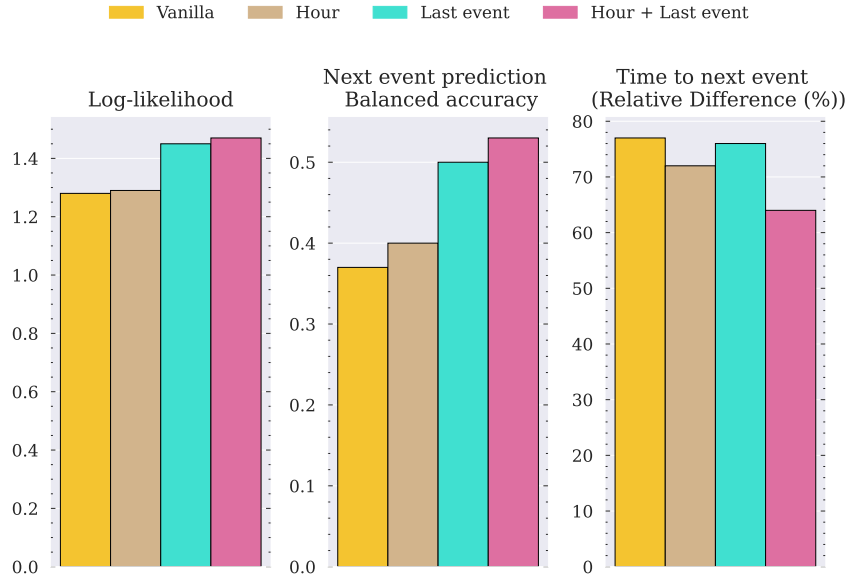


Figure 3: Model performance comparison across different feature sets: (1) Vanilla QR model ($x_k = q_k$), (2) $x_k = [q_k, h_k]$, (3) $x_k = [q_k, \eta_{k-1}]$, (4) $x_k = [q_k, h_k, \eta_{k-1}]$. Metrics: log-likelihood (higher is better), balanced accuracy of next-event prediction (higher is better), and relative difference in time to next event (lower is better).

Analysis of Results

Model Likelihood (Left Panel): Using queue size q_k alone as the baseline feature, we observe how additional features enhance the model’s performance. Adding the hour of the last event (h_k) improves the log-likelihood by capturing intra-day seasonality patterns. Incorporating the type of the last event (η_{k-1}) yields further improvement by accounting for event excitation effects. The combination of both h_k and η_{k-1} leads to the highest log-likelihood, suggesting these features provide complementary information about order flow dynamics.

Balanced Accuracy (Middle Panel): Starting from the baseline model with queue size alone, we see systematic improvements in next-event type prediction with each feature addition. While h_k captures time-of-day effects on order flow patterns, η_{k-1} shows particularly strong impact by introducing event sequencing information. Combining both features yields the highest balanced accuracy, leveraging both temporal patterns and event dependencies.

Relative Difference in Time to the Next Event (Right Panel): The baseline model provides a reference point for inter-event time prediction. Adding h_k significantly reduces the prediction error by incorporating intra-day seasonality, showing particular effectiveness in timing prediction. The addition of η_{k-1} maintains this improvement in timing accuracy, while the combined model achieves the lowest relative difference by leveraging both the temporal structure from h_k and the event sequencing from η_{k-1} .

The empirical results demonstrate the DQR framework’s capacity to capture complex market dynamics through contextual features. The enhanced performance achieved by

incorporating both h_k and η_{k-1} enables the model to reflect key market properties such as intra-day seasonality and event-type dependencies. Not only does this enriched state space improve the simulation quality, but it also enhances the model’s likelihood, validating the incorporation of additional market features. This extension of the original framework maintains a balance between model sophistication and interpretability, preserving the underlying point-process structure while accommodating richer market dynamics.

These results suggest that the DQR model’s ability to reproduce market properties depends on the features included in the state space x , with each addition potentially capturing new aspects of market behavior. This observation motivates further exploration of richer state spaces. In the next section, we present a generalization of this framework that incorporates a broader set of market features.

4. Multidimensional Deep Queue-Reactive (MDQR) Model

We have introduced the Queue-Reactive (QR) and Deep Queue-Reactive (DQR) models as frameworks for describing event arrivals in a single queue. QR models capture how different event types (limit, cancel, and market) depend on the queue size, while DQR generalizes this dependency to more complex state representations (e.g., additional market features, historical dependencies) through parametric (neural network-based) intensities. Although these models offer valuable insights, they share two critical limitations:

1. **Independence Across Queues:** Both QR and DQR consider each queue or price level in isolation. In practice, however, different price levels within a limit order book interact with one another. For example, activity at the best bid can influence behavior at the best ask, and conditions at deeper levels can propagate to the top of the book. Ignoring these cross-level dependencies limits the realism and explanatory power of the model.
2. **Absence of Order Size Modeling:** Neither QR nor DQR explicitly models the distribution of order sizes. Empirical findings [Bodor and Carlier \(2024a\)](#) indicate that order sizes play a pivotal role in shaping limit order book dynamics. By omitting order sizes, we lose an essential dimension of market complexity, potentially reducing the model’s ability to accurately reflect observed phenomena.

To address these issues, we propose the *Multidimensional Deep Queue-Reactive (MDQR)* model. The MDQR model extends the QR/DQR framework in two principal ways: (i) it treats the entire limit order book as a single multidimensional entity, thereby capturing cross-level interactions; and (ii) it incorporates order sizes into the probabilistic formulation, allowing a more comprehensive representation of market events.

Figure 4 illustrates the conceptual shift from a single-queue setting to a multidimensional one. On the left-hand side, the traditional view considers a single queue receiving three types of events (limit, cancel, market) for each queue. On the right-hand side, the MDQR model expands this perspective to multiple price levels, each capable of receiving any of the three event types. This results in a richer event space with $3 \times 2n$ distinct category-level combinations if n price levels per side are considered (where positive indices represent ask-side levels and negative indices represent bid-side levels). By unifying these

levels within a single modeling framework, the MDQR model can directly capture how conditions and events at one level influence those at another.

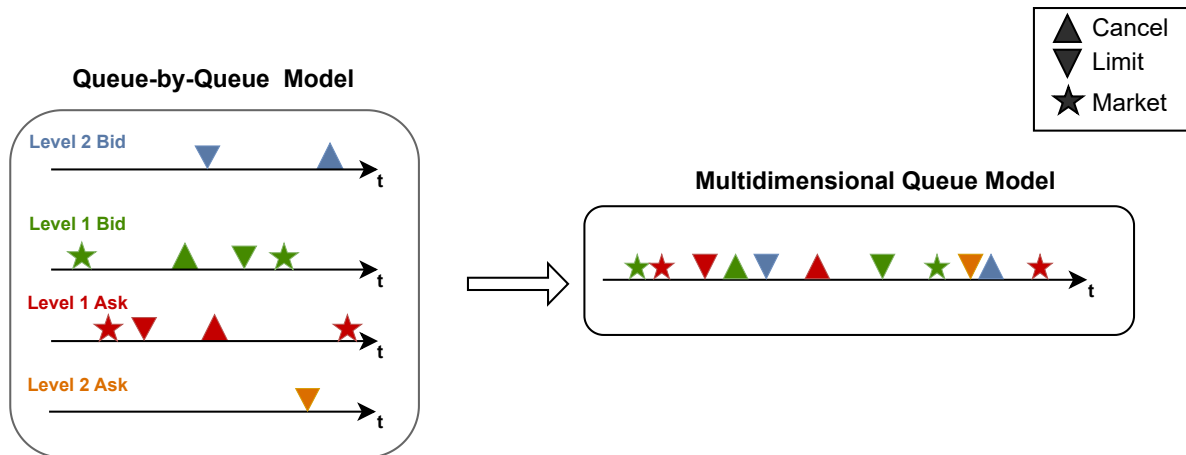


Figure 4: A conceptual illustration of the transition from a single-queue model (left) to the multidimensional MDQR framework (right). In the single-queue scenario, events occur at one level and are of three types (L, C, M). In the multidimensional scenario, multiple levels on both sides of the order book are considered together, resulting in a larger event space and enabling the model to capture cross-level interactions and dependencies.

To formalize this approach, consider a sequence of events:

$$\mathcal{E} = \{e_k\}_{k=1}^N,$$

where each event

$$e_k = (\eta_k, \ell_k, x_k, \Delta t_k, s_k)$$

is characterized by:

- $\eta_k \in \{L, C, M\}$: The event category (limit, cancel, market).
- $\ell_k \in \mathcal{P} = \{-K, \dots, -1, 1, \dots, K\}$: The price level at which the event occurs. Positive integers correspond to ask-side levels and negative integers to bid-side levels.
- x_k : A state vector that may include information from multiple levels and sides, as well as other relevant features such as historical dependencies or market indicators.
- $\Delta t_k = t_k - t_{k-1}$: The inter-arrival time since the previous event.
- s_k : The size of the order associated with the event.

The joint likelihood of the observed event sequence factorizes into two components:

$$\mathcal{L}(\theta | \mathcal{E}) = \left(\prod_{k=1}^N p(\eta_k, \ell_k, t_k | x_k) \right) \times \left(\prod_{k=1}^N p(s_k | \eta_k, \ell_k, t_k, x_k; \theta) \right) \quad (3)$$

The first part of this factorization governs the timing and category-level characteristics of events. As in the QR and DQR models, we assume a conditional Poisson structure for event arrivals. Define an intensity function $\lambda^{(\eta,\ell)}(x_k; \theta)$ for each event category η and level ℓ . The total intensity of observing any event at any level is:

$$\Lambda(x_k; \theta) = \sum_{\eta \in \{L, C, M\}} \sum_{\ell \in \mathcal{P}} \lambda^{(\eta,\ell)}(x_k; \theta)$$

The probability of an event of type η_k , at level ℓ_k occurring at time t_k given the state x_k is then:

$$p(\eta_k, \ell_k, t_k \mid x_k) \propto e^{-\Lambda(x_k; \theta) \Delta t_k} \lambda^{(\eta_k, \ell_k)}(x_k; \theta)$$

and the joint probability over all events for this part of the model is:

$$\prod_{k=1}^N p(\eta_k, \ell_k, t_k \mid x_k) \propto \prod_{k=1}^N e^{-\Lambda(x_k; \theta) \Delta t_k} \lambda^{(\eta_k, \ell_k)}(x_k; \theta).$$

This expression generalizes the DQR framework to a multidimensional setting, where the state x_k can integrate information from multiple price levels, capturing cross-level dependencies and a richer set of market signals.

For model calibration, we minimize the negative log-likelihood, which yields a more tractable optimization problem:

$$\ell_\lambda(\theta) = \sum_{k=1}^N \Lambda(x_k; \theta) \Delta t_k - \log \lambda^{(\eta_k, \ell_k)}(x_k; \theta) \quad (4)$$

where $\ell_\lambda(\theta)$ represents the objective function to be minimized to calibrate the arrival intensities model.

The second component of the likelihood,

$$\prod_{k=1}^N p(s_k \mid \eta_k, \ell_k, t_k, x_k; \theta),$$

models the distribution of order sizes given the event type, level, and state. This factor is independent of the timing and category-level intensities, allowing it to be modeled separately. In a subsequent section, we will focus on how to specify and estimate this size distribution.

In summary, the MDQR model combines a multidimensional representation of the order book with an integrated modeling of order sizes. By capturing cross-level interactions and explicitly incorporating size distributions, MDQR offers a more comprehensive and realistic framework for analyzing and simulating limit order book dynamics than previous single-queue models.

4.1. Modeling the Order Size Distribution

Recall from the MDQR model formulation that the joint likelihood of the observed event sequence factors into two components:

$$\mathcal{L}(\theta | \mathcal{E}) = \left(\prod_{k=1}^N p(\eta_k, \ell_k, t_k | x_k) \right) \times \left(\prod_{k=1}^N p(s_k | \eta_k, \ell_k, t_k, x_k; \theta) \right)$$

The first term, as discussed, governs the timing, level, and category of events and has been handled in the previous section. The second term models the distribution of order sizes conditioned on the event characteristics and state.

To specify $p(s_k | \eta_k, \ell_k, t_k, x_k; \theta)$, we require a flexible approach that can capture the diversity of size distributions observed in practice. One might consider parametric distributions such as Lognormal, Gamma, or Weibull, with parameters conditioned on (η_k, ℓ_k, x_k) . However, order size distributions often display substantial skewness and can vary dramatically depending on the event type and available liquidity. For example, sizes associated with certain event types may be effectively bounded by the remaining liquidity at that level, making it difficult for a single parametric family to fit all scenarios satisfactorily.

An alternative is to employ mixture models, such as mixtures of Gaussians. While these can approximate a wide range of distributions, they introduce additional complexity in model selection (e.g., choosing the number of components) and parameter estimation. Computational overhead and convergence stability can also become challenging, especially in settings with many conditioning variables.

Instead, a more straightforward yet effective approach is to discretize the order size space into a set of predefined classes and model the distribution as a categorical probability over these classes. This approach transforms the continuous modeling problem into a multi-class classification task, allowing us to leverage standard techniques such as cross-entropy loss minimization. The discretization can be achieved by splitting the range of order sizes into bins or quantiles, ensuring that each class corresponds to a distinct range of sizes.

In principle, if order sizes vary widely, one could define a suitable binning strategy, for example by grouping them into quantiles to maintain balanced class frequencies. This would yield a flexible, data-driven set of classes that adapt to the empirical distribution of observed order sizes.

In our specific dataset, however, over 99.9% of observed order sizes are less than 200. This natural concentration near smaller values simplifies the approach significantly. We can assign each integer size from 1 up to 200 to its own class, eliminating the need for binning or quantile-based grouping. As a result, the order size modeling reduces to predicting a categorical distribution over 200 classes (200 possible order size). A neural network can then be trained, taking as input (η_k, ℓ_k, x_k) , to produce the probability of each discrete order size class. This neural network’s parameters are estimated by minimizing the cross-entropy loss across all observed events.

Should a different dataset exhibit a larger or more dispersed range of order sizes, binning into quantiles or equal-width intervals could be employed to keep the number

of classes manageable and to maintain balanced representation across them. Thus, the approach is flexible and can be adapted to various data characteristics.

This classification-based strategy offers a balance between flexibility and practicality. It does not rely on restrictive parametric forms, avoids the complexity of mixture models, and naturally accommodates state-dependent and event-dependent patterns through the neural network’s capacity to learn non-linear relationships. This approach provides a suitable framework for modeling order sizes in the MDQR model. The likelihood maximization for this component translates to minimizing the cross-entropy loss:

$$\ell_s(\theta) = - \sum_{k=1}^N \sum_{c=1}^C y_{k,c} \log \hat{p}_c(s_k | \eta_k, \ell_k, x_k; \theta) \quad (5)$$

where \hat{p}_c denotes the model’s predicted probability for class c , $y_{k,c}$ is the one-hot encoded ground truth (1 if s_k belongs to class c , 0 otherwise), and C is the total number of size classes ($C = 200$ in our case). The probabilities are conditioned on the event characteristics (η_k, ℓ_k) and market state x_k .

4.2. Data

The data used to calibrate the MDQR model follows the framework established in Section 3.2, with one key modification in the temporal sampling approach. While the previous analysis computed Δt_k between events within individual queues, the present study considers time intervals between all order book events, providing a more comprehensive view of market dynamics. The bucketing methodology and simulation paradigm remain consistent with the earlier framework, ensuring continuity in our analytical approach.

4.3. Features

The model incorporates a diverse set of features capturing both instantaneous market state and historical market dynamics. Let us denote the limit order book state at time t_k as follows: for any level i , let $q_i(t_k)$ represent the queue size, where $i \in \{-M, \dots, -1, 1, \dots, M\}$, with negative indices corresponding to bid levels and positive indices to ask levels. Thus, $q_{-1}(t_k)$ and $q_1(t_k)$ represent the best bid and best ask queue sizes respectively.

For each price level i , we maintain a categorical variable $e_i(t_k)$ representing the type of the last event (limit order, cancellation, or market order) that occurred at that level. The spread $s(t_k)$ is computed as the difference between the best ask and best bid prices:

$$s(t_k) = p_1(t_k) - p_{-1}(t_k) \quad (6)$$

To capture order flow imbalances across different time horizons, we define the trade imbalance feature $TI_\tau(t_k)$ over a window τ as:

$$TI_\tau(t_k) = \frac{\sum_{t \in [t_k - \tau, t_k]} V^b(t) - \sum_{t \in [t_k - \tau, t_k]} V^a(t)}{\sum_{t \in [t_k - \tau, t_k]} V^b(t) + \sum_{t \in [t_k - \tau, t_k]} V^a(t)} \quad (7)$$

where $V^b(t)$ and $V^a(t)$ represent the traded volumes at the bid and ask sides respectively at time t . We consider multiple time horizons $\tau \in \{20\text{s}, 1\text{min}, 5\text{min}, 15\text{min}\}$ to capture both short-term and long-term market tendencies.

To enhance the model’s ability to learn from categorical features, each categorical variable (event types $e_i(t_k)$) is processed through an embedding layer of dimension 2, leveraging the relatively low cardinality of these features while allowing the model to learn meaningful representations of event sequences.

Table 3: Summary of Model Features and Their Characteristics

Feature Type	Description	Preprocessing	Properties
<i>Numerical Features</i>			
Queue Sizes ($q_i(t_k)$)	Level-wise volumes	Log transformation	$i \in \{-K, \dots, K\} \setminus \{0\}$
Spread ($s(t_k)$)	Price differential	None	Measured in ticks
Trade Imbalance ($TI_\tau(t_k)$)	Trades Imbalance	None	$\tau \in \{20\text{s}, 1\text{min}, 5\text{min}, 15\text{min}\}$
<i>Categorical Features</i>			
Last Event Type ($e_i(t_k)$)	Previous events	Embedding (dim=2)	Cardinality: 3
Hour ($h(t_k)$)	Current hour	Embedding (dim=2)	Cardinality: 9

4.4. MDQR model Architecture

The MDQR framework extends the architectural principles established for the DQR model (Section 3.3), adapting them to accommodate multiple price levels. The intensity prediction network employs a Multi-Layer Perceptron with two hidden layers of dimensions 256 and 64 neurons respectively. The output layer, now expanded to dimension 30, encompasses the complete event space across all monitored price levels—three possible event types (limit orders, cancellations, and market orders) for each of the ten price levels (five bid and five ask levels). The network maintains hyperbolic tangent activation functions throughout its hidden layers.

For the order size model, we preserve the core architectural elements while adjusting the network dimensions to accommodate the broader scope of market dynamics. The model retains Tanh activations in its hidden layers, with a Softmax activation in the output layer to generate appropriate probability distributions.

Table 4 presents the architectural specifications for both the intensity and order size models:

Table 4: Architectural Specifications of MDQR Components

Specification	Intensity Model	Order Size Model
Input Dimension	25	27
Hidden Layers	[256, 64]	[256, 64]
Output Dimension	30	200
Hidden Activations	Tanh	Tanh
Output Activation	ReLU	Softmax
Loss Function	Eq. (4)	Eq. (5)
Optimizer	Adam	Adam
Learning Rate Range	10^{-5} to 10^{-3}	10^{-5} to 10^{-3}

4.5. Results

To evaluate the performance of the proposed simulator, we adopt a comprehensive validation approach that examines multiple aspects of market behavior. Our analysis focuses on three key areas: the model’s ability to reproduce established market stylized facts, its capacity to capture complex patterns in order book dynamics, and its predictive power for market movements.

4.5.1. Market Impact

One of the most important use cases of limit order book (LOB) simulators is leveraging them as environments to train reinforcement learning (RL) agents or backtest execution strategies. A critical aspect of these environments is their ability to accurately replicate the market’s response to an agent’s actions, particularly concerning market impact. Market impact refers to the effect that executing a large order has on price evolution, whether the order is executed instantaneously or through a specific splitting strategy (Gatheral, 2010; Bouchaud et al., 2018).

Empirical studies have consistently observed that during a time-weighted average price (TWAP) execution of a large buy order, the price typically increases in a concave manner until reaching a peak—known as the temporary market impact—and then decreases in a convex fashion, eventually settling at a permanent market impact level. This characteristic price trajectory is depicted in Figure 5. Notable studies supporting these findings include (Zarinelli et al., 2015), (Bacry et al., 2015) and (Bouchaud et al., 2018).

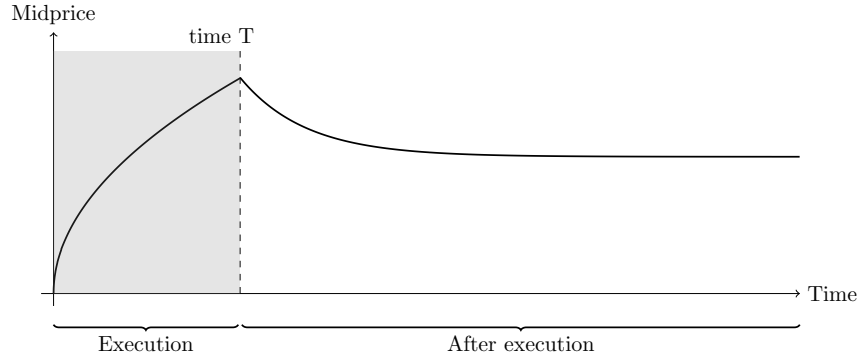


Figure 5: Average shape of price impact. During its execution, a buy metaorder drives the price upward, reaching a peak impact upon completion ($t = T$). Following the cessation of buying pressure, the price reverts sharply. However, a residual impact often remains observable well beyond the execution period and may, in some cases, persist permanently, often called permanent market impact. Figure from [Coletta et al. \(2022\)](#).

Another extensively studied phenomenon is the relationship between the magnitude of market impact and the executed quantity. The prevailing consensus is that the maximum impact does not increase linearly with the size of the executed quantity but follows a square-root law. This implies that if $I(Q)$ represents the market impact for a quantity Q , then $I(Q) \propto \sqrt{Q}$. This relationship has been documented in various studies, including ([Bucci et al., 2019](#)), ([Sato and Kanazawa, 2024](#)), and ([Gabaix et al., 2003](#)).

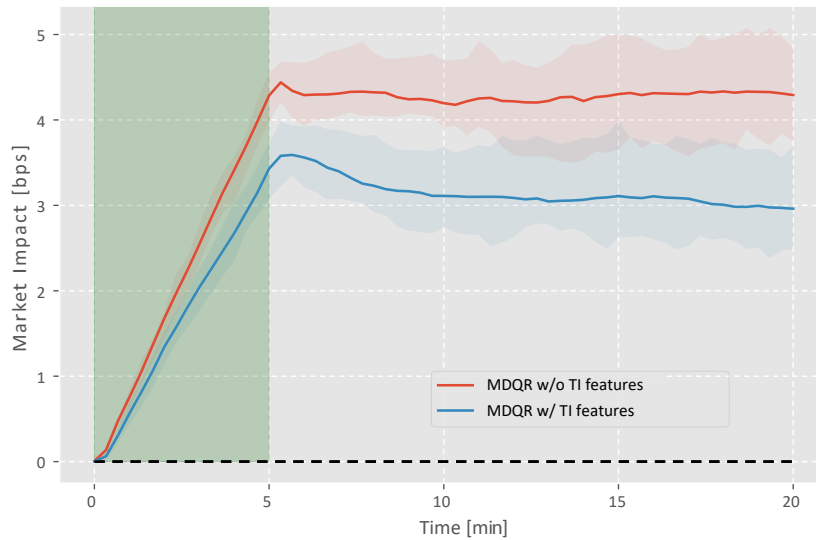


Figure 6: Market impact profiles for two models: one using only queue sizes and the other incorporating trade imbalance features. Results are based on a TWAP agent executing 100% of the average traded volume within a 5-minute window. The order is split into 10 child orders executed at a rate of 30 seconds.

The market impact profiles displayed in Figure 6 reveal distinct behavioral patterns between models with and without Trade Imbalance (TI) features. The model without TI

features exhibits a linear increase in price impact during the execution phase, followed by a plateau phase where the impact stabilizes. In contrast, the model incorporating TI features demonstrates a more nuanced and theoretically grounded behavior: it shows a concave-like progression during the execution phase, followed by a gradual relaxation period that stabilizes approximately 10 minutes after execution completion.

This characteristic shape aligns with the empirical findings of [Coletta et al. \(2022\)](#), who documented similar market impact dynamics. Furthermore, the observed pattern closely matches theoretical predictions and corresponds to the theoretical market impact profile presented in [Figure 5](#). The inclusion of trade imbalance features thus enhances the MDQR model’s ability to capture realistic market dynamics, particularly in reproducing the well-documented non-linear nature of price impact and its subsequent relaxation phase.

The improved realism in market impact modeling demonstrates that incorporating trade imbalance features allows the model to better capture the complex interplay between order flow and price formation processes in financial markets.

Furthermore, [Figure 7](#) provides additional insights into the model’s behavior under varying trading volumes. The left panel illustrates the evolution of market impact profiles for different initial quantities q_0 (quantity to buy at $T = 5$ minutes), expressed as percentages of the average 5-minute trading volume in the Bund futures market (approximately 4,000 lots, which is notably consistent with typical liquidity levels in this major fixed-income futures market). As the inventory size increases from 0% to 200% of the average traded volume in 5 minutes in the Bund market, we observe a systematic increase in market impact magnitude while maintaining similar temporal dynamics. The right panel provides empirical validation of the square-root law of market impact: plotting the maximum impact against inventory size on logarithmic scales shows a fit to a $x^{0.55}$ power law with $R^2 = 0.89$, indicating that the model reproduces this market characteristic.

The improved realism in market impact modeling demonstrates that incorporating trade imbalance features allows the model to better capture the complex interplay between order flow and price formation processes in financial markets. This advancement represents a significant step toward more accurate representation of market microstructure dynamics in quantitative trading models.




Figure 7: Average price change for different inventories (left) and maximum market impact ($I(Q)$) as a function of initial inventory size (in log-log scale) with the exponent best fit (right). Vertical lines indicate the 95% confidence intervals.

4.5.2. Excitation between different sides

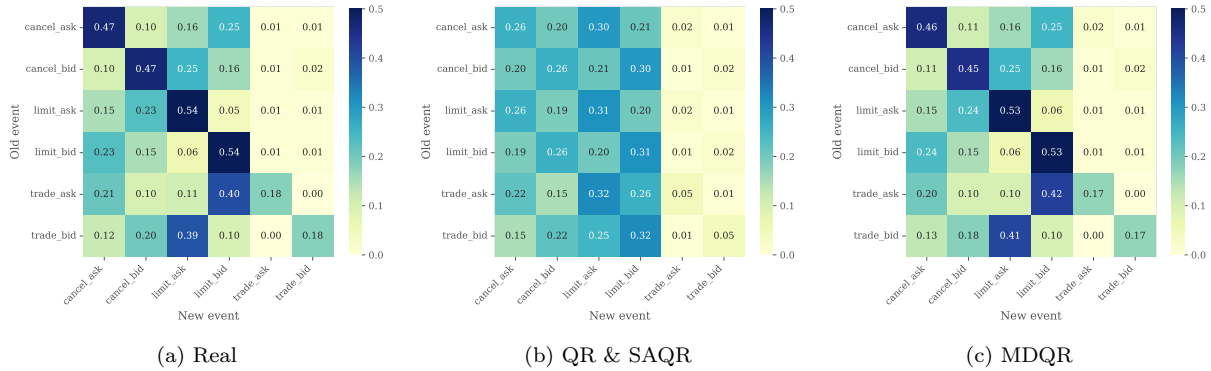


Figure 8: Transition matrix of events on simulated markets vs historical for events on the best prices.

A key aspect of market microstructure modeling is understanding the temporal dependencies between different order book events. The transition matrix $\mathbb{P}(E_{k+1}|E_k)$, where E_k represents the event type at step k , captures these dependencies by measuring the probability of transitioning from one event type to another. One of the primary motivations for relaxing the independence assumption between queues is to better model the interactions between different price levels, particularly between the best ask and best bid prices.

Figure 8 presents a comparative analysis of these transition matrices across three scenarios: historical data, SAQR model, and MDQR model. In the SAQR model, we observe that the rows of the transition matrix exhibit similar patterns, indicating minimal event-specific excitation effects. This similarity suggests that the SAQR model, with its independent queue assumption, fails to capture the complex interactions present in the order book. In contrast, the MDQR model, leveraging its ability to model different price levels simultaneously, successfully reproduces transition probabilities that closely match the historical patterns. This improved accuracy is particularly evident in the

cross-excitation effects between ask and bid events, demonstrating the model’s capability to capture the subtle interdependencies that characterize real market dynamics.

4.6. Mid-Price Prediction

While the primary objective of our framework is to generate realistic order book dynamics, evaluating its predictive capabilities provides additional validation of its ability to capture price formation mechanisms. We adopt the methodology of [Zhang et al. \(2019\)](#) to assess short-term price prediction performance, with modifications to account for the specific characteristics of the highly liquid Bund futures market.

The prediction task involves classifying directional price movements over a horizon of $k = 500$ events, where this longer horizon reflects the higher liquidity of the Bund market compared to equity markets. For each time point t_0 , we compute:

$$m_- = \frac{1}{k} \sum_{j=0}^{k-1} x_j^{\text{mid}}, \quad m_+ = \frac{1}{k} \sum_{j=1}^k x_j^{\text{mid}}, \quad r = m_+ - m_-$$

where x_j^{mid} represents the mid-price j events before t_0 for $j < 0$, at t_0 for $j = 0$, or j events after t_0 for $j > 0$. Price movements are classified as upward when $r > \Delta r$, downward when $r < -\Delta r$, and stationary otherwise. The threshold Δr is set to 20 basis points to achieve balanced class distributions in the Bund market context.

For simulator-based predictions, we generate 500 forward events starting from each observed order book state. The predicted movement is determined by computing \hat{m}_+ from these simulated trajectories and comparing $\hat{m}_+ - m_-$ against the threshold Δr .

The DeepLOB benchmark model is calibrated using one month of limit order book data, with a two-day validation period and three-day test period. The remaining data is used for training, with early stopping implemented when the validation loss shows no improvement for 10 consecutive epochs. The model is trained on the last $T = 500$ order book states, which consist of five levels of the order book on each side. A learning rate of 1×10^{-5} and a batch size of 1024 are used during training. [Table 5](#) summarizes the proportion of price movements across training, validation, and test sets, showing reasonable balance between the three classes across all datasets.

Table 5: Proportion of Price Movements in Datasets

Price Move	Train	Validation	Test
Stationary	0.38	0.31	0.37
Up	0.31	0.36	0.32
Down	0.31	0.33	0.31

Table 6: Comparison of Models: Balanced Accuracy and F1 Score.

Model	Balanced Accuracy	F1 Score
DeepLOB	0.54 ± 0.01	0.56 ± 0.01
QR	0.56 ± 0.01	0.55 ± 0.02
SAQR	0.58 ± 0.03	0.58 ± 0.03
MDQR	0.63 ± 0.02	0.62 ± 0.02

Table 6 presents the predictive performance of various models in forecasting mid-price movements. While the DeepLOB architecture, specifically designed for price prediction, achieves a baseline performance with balanced accuracy of 0.54, the generative models show competitive or superior results. The QR model matches this performance despite its simpler architecture, while the SAQR demonstrates modest improvements in both metrics. The MDQR framework achieves the highest performance across both metrics, with a balanced accuracy of 0.63 and F1 score of 0.62, suggesting that its enhanced ability to capture cross-price level dependencies contributes to improved predictive power. These results are particularly noteworthy given that price prediction is not the primary objective of these generative models.

4.7. Other Stylized Facts

4.7.1. Queue sizes distribution

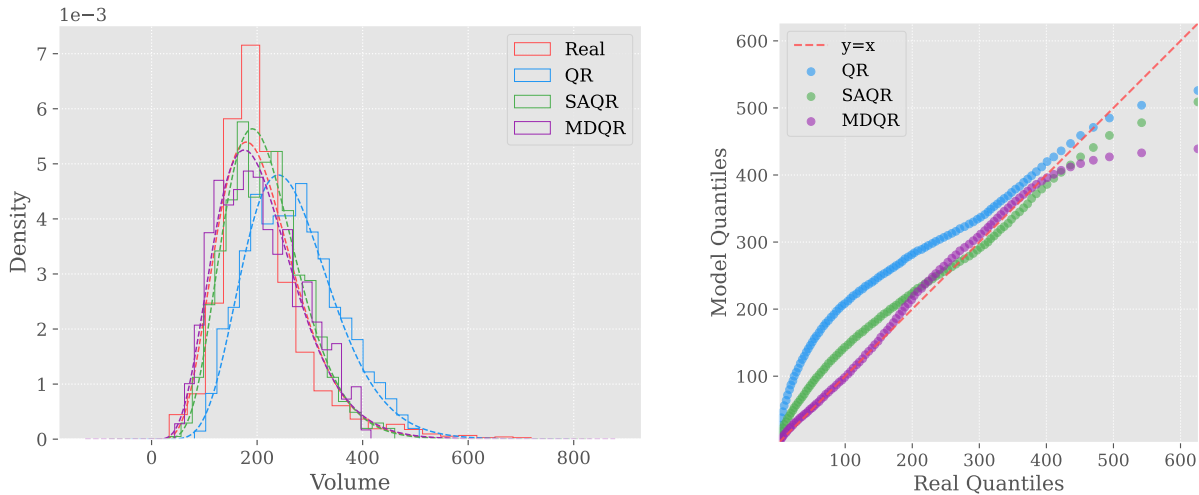


Figure 9: Comparison of real and simulated order book queue sizes. **Left:** Distribution of best ask volumes from real data versus different QR model variants, alongside a gamma distribution fit. **Right:** Q-Q plot comparing simulated market queue sizes against real queue sizes, illustrating the alignment between model predictions and actual data.

The distributional properties of queue sizes provide crucial insights into the models' ability to capture fundamental market dynamics. Figure 9 presents a comparative analysis of best ask volumes across different model specifications. The MDQR model demonstrates

Model	α	$1/\beta$
Real	1.35 ± 0.18	183.44 ± 31.79
QR	3.08 ± 0.13	83.43 ± 3.06
SAQR	1.91 ± 0.06	110.49 ± 3.14
MDQR	1.30 ± 0.13	172.00 ± 12.26

Table 7: Characteristics of gamma fit for shape and scale parameters (mean \pm std)

notable accuracy in reproducing the empirical distribution, particularly evident in the left panel where the simulated density closely tracks the historical pattern. This alignment is further validated through the Q-Q plot analysis, showing remarkable consistency between model-generated and empirical quantiles.

A particularly noteworthy aspect is the close correspondence between the empirical distribution and a gamma distribution fit, aligning with documented stylized facts in fixed income markets (Bodor and Carlier, 2024b). The gamma distribution, characterized by its probability density function:

$$f(x; \alpha, \beta) = \frac{\beta^\alpha x^{\alpha-1} e^{-\beta x}}{\Gamma(\alpha)}$$

where α represents the shape parameter and β the rate parameter (with scale = $1/\beta$), provides a natural framework for modeling queue sizes. Table 7 shows that the MDQR model closely reproduces the empirical distribution’s parameters, with shape ($\alpha = 1.30 \pm 0.13$) and scale ($1/\beta = 172.00 \pm 12.26$) values closely matching those of historical data ($\alpha = 1.35 \pm 0.18$, $1/\beta = 183.44 \pm 31.79$). This alignment suggests that our model effectively captures the underlying properties of queue size formation, while simpler models like QR and SAQR show notable deviations in their distributional parameters.

Figure 10 extends this analysis to deeper price levels, comparing average posted volumes between models and historical data. The MDQR framework maintains its superior performance across different price levels, consistently producing volume patterns that closely match empirical observations. This broader success in reproducing volume characteristics suggests that the model effectively captures not just the dynamics at the best quotes, but also the complex interactions that shape liquidity provision across the order book’s depth.

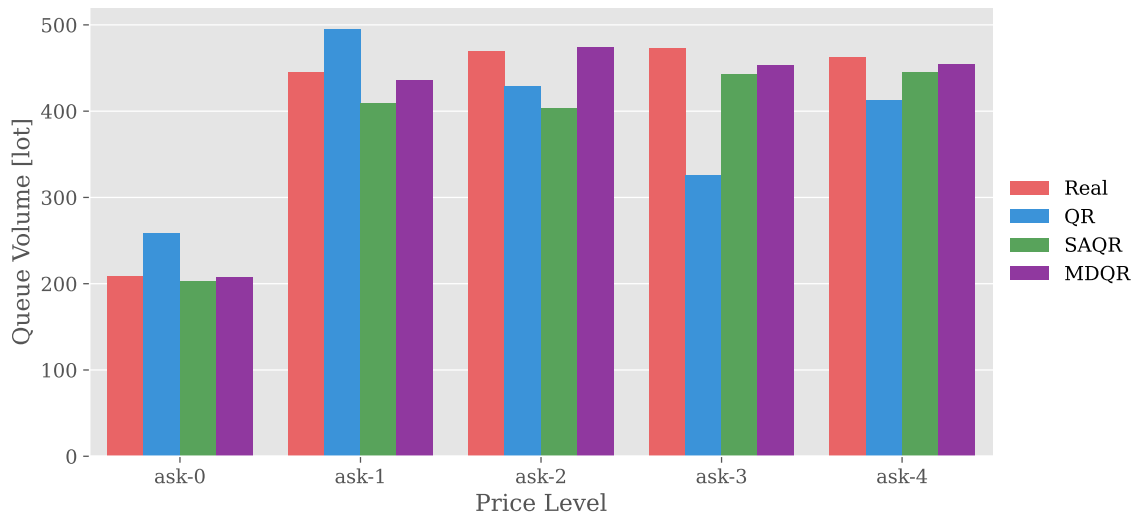


Figure 10: Average LOB volumes for the first five levels at the ask side.

4.7.2. Correlation between Order Book Queues

The correlation structure between order book queues provides crucial insights into the models' ability to capture complex market microstructure dynamics. Figure 11 presents correlation matrices for queue sizes across different price levels, comparing empirical patterns with those generated by the QR, SAQR, and MDQR models.

The empirical correlation matrix (Figure 11a) reveals several notable patterns. First, there exists a strong negative correlation (-0.54) between the best bid and ask queues, reflecting the natural asymmetry in order flow: periods of high selling pressure often coincide with depleted buying interest and vice versa. Second, we observe significant positive correlations between queues on the same side of the book, with values typically ranging from 0.3 to 0.5, indicating coordinated liquidity provision behaviors.

The QR model (Figure 11b) fails to capture these essential relationships. Most strikingly, it underestimates the negative correlation between best bid and ask (-0.22 versus -0.54 in empirical data) and shows negligible correlations between queues at different price levels, with most values close to zero. This limitation stems from the model's independent treatment of individual queues.

While the SAQR model (Figure 11c) shows modest improvement in capturing the best bid-ask correlation (-0.43), it still struggles to reproduce the broader correlation structure, particularly the positive correlations between queues on the same side of the book. Most off-diagonal elements remain close to zero, indicating that the model fails to capture the interconnected nature of liquidity dynamics.

The MDQR model (Figure 11d) demonstrates markedly superior performance in reproducing the empirical correlation structure. It accurately captures both the negative correlation between best bid and ask (-0.54) and the positive correlations between same-side queues. The model generates realistic correlation patterns that extend beyond immediate price levels, suggesting its effectiveness in capturing the complex interplay of liquidity provision and consumption across the order book, all of this thanks to the MDQR ability

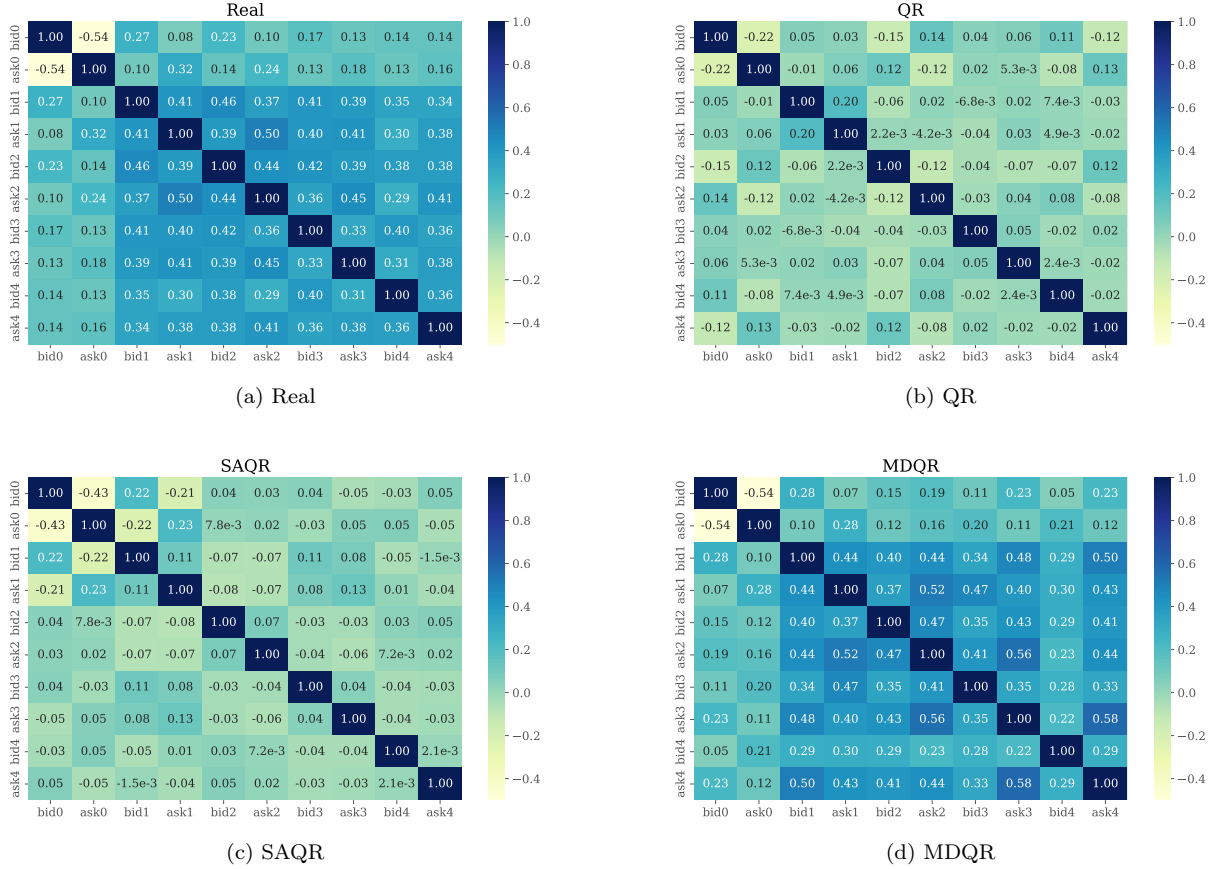


Figure 11: Correlation matrices of queue volumes across different price levels (5 bid and 5 ask levels) for historical data (Real) and three model variants (QR, SAQR, and MDQR). Each entry (i, j) represents the Pearson correlation coefficient between volumes at price levels i and j , where negative indices denote bid levels and positive indices denote ask levels.

to take into account states from multiples queues to simulate events.

4.7.3. Distribution of Returns

The distributional properties of one-minute returns provide valuable insights into the models' capacity to capture market dynamics. Figure 12 presents both the probability density functions and quantile-quantile plots for the various model specifications. The MDQR model demonstrates superior performance in reproducing the empirical distribution of returns, particularly in capturing the characteristic heavy tails of the historical returns.

This improved distributional fit is further validated through the Q-Q plot analysis, where the MDQR quantiles closely align with the theoretical line $y = x$, indicating strong agreement between simulated and empirical distributions across the entire range of returns. In contrast, the QR model exhibits notable deviations, particularly in the tail regions. As discussed in Bodor and Carlier (2024a), these limitations stem from the QR model's inability to adequately capture rapid price movements due to its simplified queue dynamics.

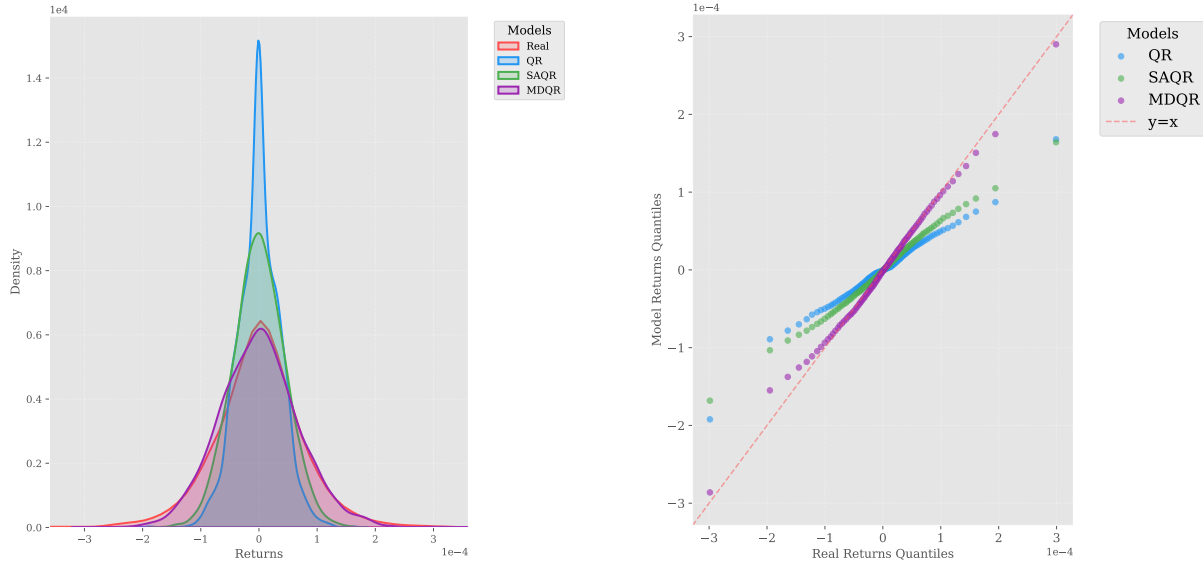


Figure 12: **Left:** Distribution of one-minute returns for real data and various model variants, highlighting differences in return profiles. **Right:** Q-Q plot comparing model-generated returns with real returns, showing the degree of alignment between simulated and actual data.

The enhanced distributional accuracy of the MDQR model can be attributed to its more sophisticated treatment of cross-price level interactions and the incorporation of market-wide state variables, allowing it to better reflect the complex mechanisms driving price formation in real markets.

4.7.4. Frequency of events

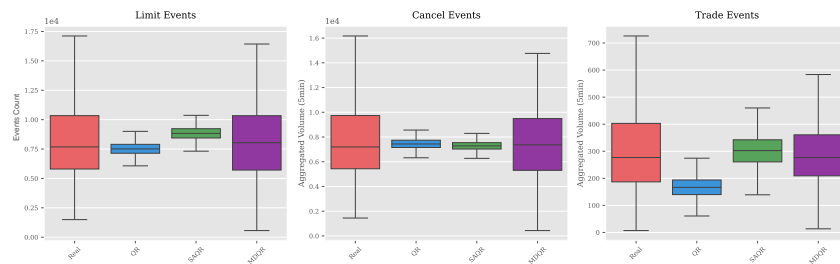


Figure 13: Distribution of event counts per 5-minute window across different model specifications. Box plots show the comparison between historical data (Real) and three model variants (QR, SAQR, and MDQR) for limit orders, cancellations, and trades.

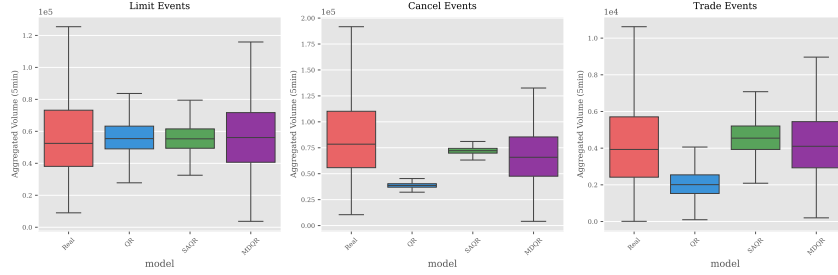


Figure 14: Distribution of aggregated volumes per 5-minute window across different model specifications. Box plots compare historical data (Real) with QR, SAQR, and MDQR models for limit orders, cancellations, and trades, showing both median values and dispersion patterns.

Figures 13 and 14 present a comprehensive comparison of event distributions across different model specifications through box plots, examining both event counts and aggregated volumes within 5-minute windows. The analysis reveals several key insights about the models’ capabilities and limitations.

The QR model demonstrates accurate reproduction of average behavior for limit and cancel orders, both in terms of event counts and volumes. However, it shows significant discrepancy in modeling trade events. The SAQR extension successfully addresses this limitation, bringing the trade event distributions closer to historical patterns. This improvement highlights the importance of incorporating additional queue dependencies in the modeling framework.

Nevertheless, both QR and SAQR models, relying solely on queue sizes, exhibit a notable limitation: they fail to capture the true market variability. The observed standard deviations in these models are considerably smaller than those present in historical data, indicating an oversimplified representation of market dynamics.

The MDQR model addresses these shortcomings by not only reproducing average behaviors accurately but also capturing a wider range of market scenarios. The distributions more closely match the historical data across all event types, particularly in terms of variability. This enhanced capability stems from the model’s ability to incorporate multiple price levels and their interactions, resulting in a more realistic representation of market microstructure dynamics.

4.7.5. Simulation Time

Model	Inference Time (ms)	Daily Generation (min)
MDQR	0.084 ± 0.007	2.10
LOGAN (Sirignano and Cont, 2019)	0.402 ± 0.081	10.05
RNN (Hultin et al., 2023)	1.152 ± 0.061	28.85
WGAN (Cont et al., 2023)	0.144 ± 0.008	3.60

Table 8: Performance comparison of different architectures for limit order book modeling. Inference time is reported as mean \pm standard deviation for generating a single event. Daily generation time represents the total computation time required to generate 1.5M events, equivalent, on average, to one trading day (9 AM to 6 PM) in the Bund Market.

The computational efficiency of order book simulators plays a crucial role in various applications. While generating synthetic data might require a reasonable number of simulations, making computational speed less critical, reinforcement learning algorithms typically need billions of market events for each training run. Moreover, these simulations generally cannot be reused between different training iterations. This makes computational efficiency a crucial factor for practical implementation.

Table 8 presents a performance comparison across different limit order book modeling approaches. The MDQR framework demonstrates remarkable computational efficiency, achieving an inference time² of 0.084 milliseconds per event. This performance translates to just 2.10 minutes for simulating an entire trading day’s worth of events³, representing a significant advantage over existing approaches. In contrast, while architectures like LOBGAN Sirignano and Cont (2019) and RNN Hultin et al. (2023) may offer sophisticated feature extraction capabilities, they require substantially more computational resources, with daily generation times ranging from 10 to 29 minutes.

These performance considerations motivated our choice of a simple MLP architecture over potentially more sophisticated approaches. While advanced architectures might offer enhanced feature extraction capabilities, the specific requirements of reinforcement learning applications necessitate finding an optimal balance between model complexity and computational efficiency. This underscores the importance of considering the intended use case when selecting an appropriate modeling architecture.

4.7.6. Order Sizes Distribution

A key innovation of the MDQR framework lies in its treatment of order size modeling. Unlike the SAQR model, which relies on independent sampling mechanisms, MDQR implements a conditional approach that considers the broader market context when generating order sizes. This methodology allows the model to capture the subtle interplay between market conditions and traders’ size choices, reflecting how market participants adapt their order sizes based on prevailing market conditions and order book state.

²Add Cpu details

³Reported times account only for event generation and do not include the computational overhead of updating the order book through the matching engine, which is common to all approaches.




Figure 15: Fitted distribution (left) and the corresponding QQ-plot (right) for trade orders arriving when the size of the queue is 15.




Figure 16: Fitted distribution (left) and the corresponding QQ-plot (right) for limit orders arriving when the size of the queue is 20.

The distribution of order sizes exhibits distinct patterns depending on both the order type and market conditions. Figures 15 and 16 illustrate this complexity through two representative cases. Figure 15 shows the distribution of trade order sizes when the queue size is 15 lots, revealing a specific pattern with notable peaks at certain size values (especially 1 lot and 15 lots). In contrast, Figure 16 presents the distribution of limit order sizes for a queue size of 20 lots, displaying a different structure with higher concentration at smaller sizes. The MDQR sizing model appears to capture these distinct conditional distributions, as evidenced by the close alignment between empirical and simulated patterns in both cases. The corresponding Q-Q plots support this observation, showing consistent agreement across the range of order sizes for both trade and limit orders.

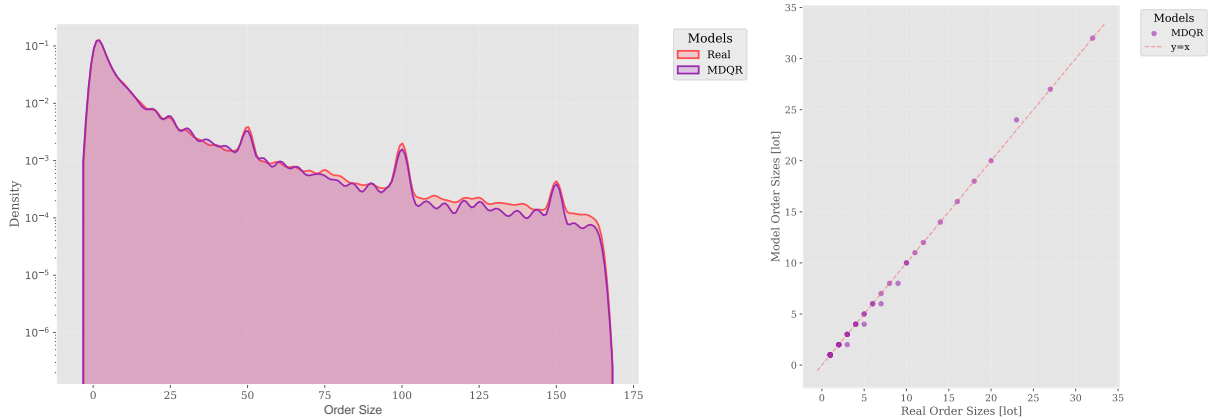


Figure 17: **Left:** Stationary distribution of order sizes for real data and the MDQR model, illustrating the similarity between the two distributions. **Right:** Q-Q plot comparing the stationary distribution of order sizes generated by the MDQR model with real data, demonstrating the level of agreement between the simulated and actual distributions.

Figure 17 presents evidence of the model’s effectiveness in capturing overall order size dynamics. The left panel displays the stationary distribution of order sizes, where the MDQR-generated distribution mirrors the historical pattern across the entire range of observed sizes. This macroscopic alignment suggests that the model internalizes the factors influencing traders’ size decisions. The Q-Q plot in the right panel provides quantitative validation of this distributional accuracy, showing alignment between simulated and empirical quantiles.

These results indicate that the MDQR sizing component can reproduce order size distributions at different scales. The model appears to capture both the specific conditional distributions that arise under particular market conditions and the overall stationary distribution of sizes, suggesting its utility for realistic market simulation.

5. Comparison with Recent Deep Learning Approaches

Our deep learning generalization of the QR model contributes to a growing body of research leveraging neural networks for limit order book simulation. Notable recent approaches include the conditional GAN framework of [Coletta et al. \(2022\)](#), the Wasserstein GAN with gradient penalty proposed by [Cont et al. \(2023\)](#), and the multi-model RNN architecture developed by [Hultin et al. \(2023\)](#), which employs separate networks for event type, price level, size, and arrival time generation.

These approaches demonstrate remarkable capability in reproducing market stylized facts. [Coletta et al. \(2022\)](#) successfully captures the characteristic “concave-convex” shape of large order execution impact, a feature also observed in [Hultin et al. \(2023\)](#) and our MDQR model. While [Cont et al. \(2023\)](#) does not fully capture the relaxation phase post-execution, it effectively reproduces the square-root law of market impact, demonstrating that maximum impact scales with the square root of executed position size—a property our model also verifies.

In terms of order book structure, [Hultin et al. \(2023\)](#) and MDQR both accurately reproduce the average shape across different price levels, while [Cont et al. \(2023\)](#) extends this to capture the complete distribution of queue sizes. Additionally, our MDQR simulator demonstrates strong performance in short-horizon mid-price movement prediction, achieving comparable or superior results to specialized prediction models, aligning with similar findings by [Hultin et al. \(2023\)](#).

While these approaches share common objectives, they offer complementary strengths. The GAN-based models of [Coletta et al. \(2022\)](#) and [Cont et al. \(2023\)](#), which are architecturally designed for generative tasks, face training stability challenges requiring careful calibration. These models generate order book states at fixed intervals (e.g., 10-second snapshots in [Cont et al. \(2023\)](#)), which proves advantageous when high-frequency precision is unnecessary but limits access to detailed order flow information between snapshots, which limits usage of such information in the modeling.

In contrast, [Hultin et al. \(2023\)](#) and MDQR model event-by-event arrival dynamics. The MDQR framework, while not yet validated for small-tick assets, contrasts with [Hultin et al. \(2023\)](#) and [Coletta et al. \(2022\)](#), which demonstrate effectiveness in reproducing spread distributions and regime transitions in such markets. Furthermore, [Coletta et al. \(2022\)](#), [Cont et al. \(2023\)](#), and MDQR employ manually crafted state spaces, while [Hultin et al. \(2023\)](#) leverages automatic feature extraction through LSTM architecture. This distinction presents a trade-off between model complexity control and feature engineering requirements—our implementation, for instance, required several iterations before incorporating trade imbalance features crucial for accurate market impact profiles.

Computational efficiency varies significantly across approaches. MDQR achieves superior inference speed due to its streamlined architecture, while [Hultin et al. \(2023\)](#)’s four-model structure results in longer computation times. Our model provides an interpretable framework based on Poisson event arrival while maintaining the ability to reproduce key market stylized facts with efficient simulation capabilities.

The selection of an appropriate model depends on specific use cases. MDQR proves suitable for high-frequency simulation of liquid assets, particularly when computational efficiency is most important. [Cont et al. \(2023\)](#)’s approach may be preferable for lower-frequency applications, despite potential limitations in order flow information. For small-tick assets, [Cont et al. \(2023\)](#) and [Hultin et al. \(2023\)](#) demonstrate particular effectiveness. When automatic pattern extraction takes precedence over simulation speed, [Hultin et al. \(2023\)](#)’s approach offers distinct advantages through its LSTM architecture.

6. Conclusion

This paper presents the Multidimensional Deep Queue-Reactive (MDQR) model, developed through a systematic enhancement of the QR framework. Our initial investigation with the DQR model demonstrated how deep learning techniques could effectively generalize the state space of queue-reactive models, with each additional feature capturing new market properties in the simulated dynamics. Building on these insights, we developed the MDQR framework, which further extends this approach by relaxing the queue independence assumption and introducing a more sophisticated treatment of order sizes.

Our analysis using Bund futures market data suggests the model can capture various market characteristics, including market impact profiles, event-type dependencies across order book sides, and mid-price movement patterns, among others. The MDQR framework incorporates multiple price levels and market state features while retaining computational efficiency. This approach allows for the inclusion of relevant market features in the state space, helping to reproduce observed market behavior in simulations. By maintaining its connection to point processes, the model provides insights into the relationships between market features and order flow patterns. This work contributes to the development of market simulation tools by combining elements from queue-reactive models and deep learning methods. The resulting framework aims to balance model sophistication with practical implementation, while remaining adaptable to evolving market dynamics.

The current implementation shows promising results for liquid and large-tick assets such as the Bund futures and opens several directions for future research. Testing the model’s applicability to small-tick assets with different market characteristics would help validate its broader utility. Additionally, the model’s computational efficiency makes it particularly suitable for reinforcement learning applications in finance. Future work will focus on leveraging this framework for developing automated trading strategies, particularly in areas such as market making and optimal execution.

7. Ablation Study

Here, I plan to add an ablation study about the impact of different state space features on the intensities and arrival of orders.

For example, how the ratio ‘ask trade rate’ / ‘bid trade rate’ evolve with OB imbalance ? How limit order arrival (at each side) behave w.r.t trade imbalance ? How the excitation between events evolve during the day ? during order book imbalance phases ? and other questions like this.

References

- F. Abergel, A. Jedidi, A mathematical approach to order book modeling, *International Journal of Theoretical and Applied Finance* 16 (05) (2013) 1350025.
- F. Abergel, A. Jedidi, Long-time behavior of a Hawkes process-based limit order book, *SIAM Journal on Financial Mathematics* 6 (1) (2015) 1026–1043.
- R. L. Axtell, J. D. Farmer, Agent-based modeling in economics and finance: Past, present, and future, *Journal of Economic Literature* (2022) 1–101.
- E. Bacry, A. Iuga, M. Lasnier, C.-A. Lehalle, Market impacts and the life cycle of investors orders, *Market Microstructure and Liquidity* 1 (02) (2015) 1550009.
- H. Bodor, L. Carlier, A Novel Approach to Queue-Reactive Models: The Importance of Order Sizes, arXiv preprint arXiv:2405.18594 .

- H. Bodor, L. Carlier, Stylized Facts and Market Microstructure: An In-Depth Exploration of German Bond Futures Market, arXiv preprint arXiv:2401.10722 .
- J.-P. Bouchaud, J. Bonart, J. Donier, M. Gould, Trades, quotes and prices: financial markets under the microscope, Cambridge University Press, 2018.
- J.-P. Bouchaud, M. Mézard, M. Potters, et al., Statistical properties of stock order books: empirical results and models, *Quantitative Finance* 2 (4) (2002) 251–256.
- F. Bucci, M. Benzaquen, F. Lillo, J.-P. Bouchaud, Crossover from linear to square-root market impact, *Physical review letters* 122 (10) (2019) 108302.
- Á. Cartea, S. Jaimungal, J. Penalva, Algorithmic and high-frequency trading, Cambridge University Press, 2015.
- A. Coletta, et al., Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness, arXiv preprint arXiv:2206.08351 .
- R. Cont, Statistical modeling of high-frequency financial data, *IEEE Signal Processing Magazine* 28 (5) (2011) 16–25.
- R. Cont, M. Cucuringu, J. Kochems, F. Prezel, Limit Order Book Simulation with Generative Adversarial Networks, arXiv preprint arXiv:2306.12806 .
- R. Cont, A. De Larrard, Price dynamics in a Markovian limit order market, *SIAM Journal on Financial Mathematics* 4 (1) (2014) 1–25.
- R. Cont, M. S. Müller, A stochastic partial differential equation model for limit order book dynamics, *SIAM Journal on Financial Mathematics* 12 (2) (2021) 744–787.
- R. Cont, S. Stoikov, R. Talreja, A stochastic model for order book dynamics, *Operations Research* 58 (3) (2010) 549–563.
- X. Gabaix, P. Gopikrishnan, V. Plerou, H. E. Stanley, A theory of power-law distributions in financial market fluctuations, *Nature* 423 (6937) (2003) 267–270.
- J. Gatheral, No-dynamic-arbitrage and market impact, *Quantitative Finance* 10 (7) (2010) 749–759.
- M. D. Gould, M. A. Porter, S. Williams, M. McDonald, D. J. Fenn, S. D. Howison, Limit order books, *Quantitative Finance* 13 (11) (2013) 1709–1742.
- W. Huang, C.-A. Lehalle, M. Rosenbaum, Simulating and analyzing order book data: The queue-reactive model, *Journal of the American Statistical Association* 110 (509) (2015) 107–122.
- H. Hultin, et al., A generative model of a limit order book using recurrent neural networks, *Quantitative Finance* 23 (6) (2023) 931–950.

- A. Kirilenko, A. S. Kyle, M. Samadi, T. Tuzun, The flash crash: High-frequency trading in an electronic market, *The Journal of Finance* 72 (3) (2017) 967–998.
- M. Paddrik, R. Hayes, A. Todd, S. Yang, P. Beling, W. Scherer, An agent based model of the E-Mini S&P 500 applied to flash crash analysis, 2012 IEEE Conference on Computational Intelligence for Financial Engineering & Economics (CIFER) (2012) 1–8.
- M. Potters, J.-P. Bouchaud, More statistical properties of order books and price impact, *Physica A: Statistical Mechanics and its Applications* 324 (1-2) (2003) 133–140.
- Y. Sato, K. Kanazawa, Does the square-root price impact law belong to the strict universal scalings?: quantitative support by a complete survey of the Tokyo stock exchange market, arXiv preprint arXiv:2411.13965 .
- J. Sirignano, R. Cont, Universal features of price formation in financial markets: perspectives from deep learning, *Quantitative Finance* 19 (9) (2019) 1449–1459.
- E. Smith, J. D. Farmer, L. s. Gillemot, S. Krishnamurthy, et al., Statistical theory of the continuous double auction, *Quantitative Finance* 3 (6) (2003) 481–514.
- S. Vyetrenko, D. Byrd, N. Petosa, M. Mahfouz, D. Dervovic, M. Veloso, T. Balch, Get real: Realism metrics for robust limit order book market simulations, in: *Proceedings of the First ACM International Conference on AI in Finance*, 1–8, 2020.
- P. Wu, M. Rambaldi, J.-F. Muzy, E. Bacry, Queue-reactive Hawkes models for the order flow, arXiv preprint arXiv:1901.08938 .
- E. Zarinelli, M. Treccani, J. D. Farmer, F. Lillo, Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate, *Market Microstructure and Liquidity* 1 (02) (2015) 1550004.
- Z. Zhang, S. Zohren, S. Roberts, Deeplob: Deep convolutional neural networks for limit order books, *IEEE Transactions on Signal Processing* 67 (11) (2019) 3001–3012.

Appendix A. Complementary Stylized Facts Results

Appendix A.1. Appendix section