

DiffLOB: Diffusion Models for Counterfactual Generation in Limit Order Books

Zhuohan Wang, Carmine Ventre

King’s College London

{zhuohan.wang, carmine.ventre}@kcl.ac.uk

Abstract

Modern generative models for limit order books (LOBs) can reproduce realistic market dynamics, but remain fundamentally passive: they either model what typically happens without accounting for hypothetical future market conditions, or they require interaction with another agent to explore alternative outcomes. This limits their usefulness for stress testing, scenario analysis, and decision-making. We propose **DiffLOB**, a regime-conditioned **Diffusion** model for controllable and counterfactual generation of **LOB** trajectories. DiffLOB explicitly conditions the generative process on future market regimes—including trend, volatility, liquidity, and order-flow imbalance, which enables the model to answer counterfactual queries of the form: “If the future market regime were X instead of Y, how would the limit order book evolve?” Our systematic evaluation framework for counterfactual LOB generation consists of three criteria: (1) *Controllable Realism*, measuring how well generated trajectories can reproduce marginal distributions, temporal dependence structure and regime variables; (2) *Counterfactual validity*, testing whether interventions on future regimes induce consistent changes in the generated LOB dynamics; (3) *Counterfactual usefulness*, assessing whether synthetic counterfactual trajectories improve downstream prediction of future market regimes.

1 Introduction

Modern financial markets are driven by automated trading systems operating on limit order books (LOBs) [Gould *et al.*, 2013]. A LOB can be thought of as two priority queues collecting buy and sell orders of market participants. Each order is comprised of bid/ask price (how much one is willing to pay/get for each stock) and volume (how many stocks one wants to trade). Small changes in trend, liquidity, volatility, or order flow balance of the LOB can lead to drastically different market outcomes. As a result, stress testing, strategy design, and market surveillance all require not only accurate models of how markets behave on average, but also the ability to simulate how the market would become under hypothetical

future conditions, such as fast trend change, liquidity shocks, volatility spikes, or imbalanced order flow.

However, existing approaches to LOB modeling exhibit fundamental limitations in answering counterfactual questions. Supervised LOB predictors typically focus on forecasting specific future variables, such as mid-price returns [Briola *et al.*, 2025a; Briola *et al.*, 2025b], but they do not generate full market trajectories and therefore cannot simulate how the order book would evolve over time under alternative future conditions. Recent generative models for LOBs are capable of producing highly realistic samples when conditioned on historical observations [Nagy *et al.*, 2023; Coletta *et al.*, 2021; Li *et al.*, 2025]. However, to obtain counterfactual outcomes, these models usually rely on explicit interactions with trading agents. As a result, the counterfactual trajectories generated are not guaranteed as a desired future state of the market.

To address this gap, we propose DiffLOB¹, a diffusion model-based framework specifically designed for the counterfactual generation of LOB snapshots. The proposed framework offers three main contributions:

- (1) We formulate counterfactual generation of LOBs as a conditional generative modeling problem and explicitly introduce future market regimes as control variables to enable controllable and counterfactual generation. We develop a novel diffusion-based architecture that conditions on past LOB trajectories, time-of-day information, and four future regime variables—trend, volatility, liquidity, and order-flow imbalance. The proposed model effectively captures key financial statistics and consistently outperforms strong baseline models across a range of evaluation metrics.
- (2) We empirically demonstrate that interventions on future market regimes induce consistent and interpretable changes in the generated LOB dynamics. Under extreme and hypothetical regime conditions, DiffLOB produces trajectories whose statistical properties align with the imposed regimes and closely match real markets observed under comparable conditions.
- (3) We show that synthetic counterfactual trajectories generated by DiffLOB provide meaningful additional infor-

¹Our code is available at <https://github.com/ZhuoHan1998/DiffLOB>.

mation for downstream tasks. In particular, augmenting real data with counterfactual samples consistently improves the performance of models for future market regime prediction under extreme conditions, demonstrating the practical value of counterfactual LOB generation.

2 Related Work

Generative Modelling in LOBs. A growing body of work applies generative models to LOBs with the goal of reproducing realistic market microstructure dynamics. Autoregressive approaches have been used to model LOB dynamics at the event level. Hultin *et al.* decompose the joint distribution of LOB transitions into conditional components using recurrent neural networks, while Nagy *et al.* propose an end-to-end autoregressive model based on structured state-space models that tokenizes message streams. GAN-based approaches, on the other hand, employ adversarial training to synthesize LOB data with high visual and statistical realism. Coletta *et al.* propose a Conditional GAN framework that reacts to current market states and allows agent interaction within a simulation, showing enhanced realism and responsiveness. Recent work has also explored diffusion-based approaches for limit order book simulation. TRADES [Berti *et al.*, 2025] develops a diffusion-based market simulator that generates order-level LOB dynamics and supports counterfactual analysis through interaction with trading agents. DiffVolume [Wang and Ventre, 2025] apply diffusion models to generate high-dimensional LOB volume snapshots across multiple price levels without capturing full LOB trajectories. Li *et al.* create an order-level generative foundation model for downstream forecasting, risk detection and financial analysis. Despite their success in realism, existing generative LOB models primarily learn the observational distribution of market dynamics. They are not designed to explicitly intervene on future market regimes, and therefore cannot generate counterfactual and complete LOB trajectories corresponding to hypothetical future conditions.

Diffusion Models for Financial Time Series. Diffusion models are a class of generative models inspired by thermodynamic diffusion processes, where data are progressively perturbed with Gaussian noise and generated by learning to reverse this process [Sohl-Dickstein *et al.*, 2015; Ho *et al.*, 2020]. This formulation is closely connected to score-based generative modeling, which learns the gradient of the data density via denoising score matching [Vincent, 2011; Song and Ermon, 2019], and has been unified under a stochastic differential equation (SDE) framework [Song *et al.*, 2021]. Building on these advances, diffusion models have been increasingly applied to financial time series. Koa *et al.* employ diffusion models for multi-step stock price prediction, while Wang and Ventre apply diffusion-based denoising techniques to financial time series. More recently, Tanaka *et al.* propose a controllable diffusion framework for financial time series generation and introduce normalization-based techniques to enhance controllability [Hashimoto *et al.*, 2025].

3 Methodology

3.1 Problem Formulation

Let $x_t \in \mathbb{R}^{K \times C}$ denote the state of LOB at time t , where K denotes the number of price levels (e.g., 10 ask price levels and 10 bid price levels) and C represents the feature dimension (e.g., price and volume). Given a historical LOB trajectory $x_{1:t}$, our goal is to generate a future trajectory $x_{t+1:t+\tau}$. We introduce future market regimes as controlling variables to enable controllable and counterfactual generation. Specifically, we define a set of future regime variables

$$c_{t+1:t+\tau} = (c_{t+1:t+\tau}^{\text{trend}}, c_{t+1:t+\tau}^{\text{vol}}, c_{t+1:t+\tau}^{\text{liq}}, c_{t+1:t+\tau}^{\text{imb}}),$$

where $\{c_{t+1:t+\tau}^{\text{trend}}, c_{t+1:t+\tau}^{\text{vol}}\} \in \mathbb{R}$, $\{c_{t+1:t+\tau}^{\text{liq}}, c_{t+1:t+\tau}^{\text{imb}}\} \in \mathbb{R}^\tau$

corresponding to trend, volatility, liquidity, and order-flow imbalance, respectively. These variables characterize the macroscopic state of the market over the future horizon and are treated as *intervenable control variables* rather than observed labels. Our objective is to model the distribution

$$p(x_{t+1:t+\tau} \mid x_{1:t}, c_{t+1:t+\tau}), \quad (1)$$

which enables counterfactual queries of the form: “How would the limit order book evolve if the future market regime conditions were different?” By explicitly conditioning on future regimes, DiffLOB decouples market dynamics from specific agent behaviors and provides direct control over hypothetical future market conditions.

We adopt a denoising diffusion probabilistic model (DDPM) [Ho *et al.*, 2020] to represent the conditional distribution of future LOB trajectories. In the forward process, Gaussian noise is progressively added to the target trajectory $x_{t+1:t+\tau}$ over a sequence of diffusion steps. A neural network is trained to reverse this process by predicting the injected noise at each step, thereby learning to generate samples from the data distribution. Conditions on historical trajectories and future market regimes are incorporated through the denoising network, allowing the diffusion process to generate future LOB trajectories consistent with specified regimes.

3.2 DiffLOB Architecture

An overview of DiffLOB architecture is shown in Figure 1.

Overall Backbone. As illustrated in Figure 1(a), DiffLOB adopts stacked Wavenet-style residual blocks [van den Oord *et al.*, 2016]. The backbone models the structural and temporal dynamics of LOB price and volume evolution, producing residual connection Res_i and skip connection $Skip_i$ at each block, where residual outputs are propagated to subsequent blocks and skip outputs are aggregated to form the final prediction of score s_θ . The input $\tilde{x}_{t+1:t+\tau}$ denotes the noised LOB trajectory at a given diffusion step. Dim_{emb} is price level embedding, which enables the model to distinguish between spatially different price levels in the order book.

Regime Condition Encoders. Future market regimes are encoded through a condition encoder (Figure 1(b)). We distinguish between *Local* and *Global* regime information. Local conditions capture step-wise signals, including historical LOB, liquidity, imbalance, and time-of-day information, and are encoded using convolutional layers. Global conditions

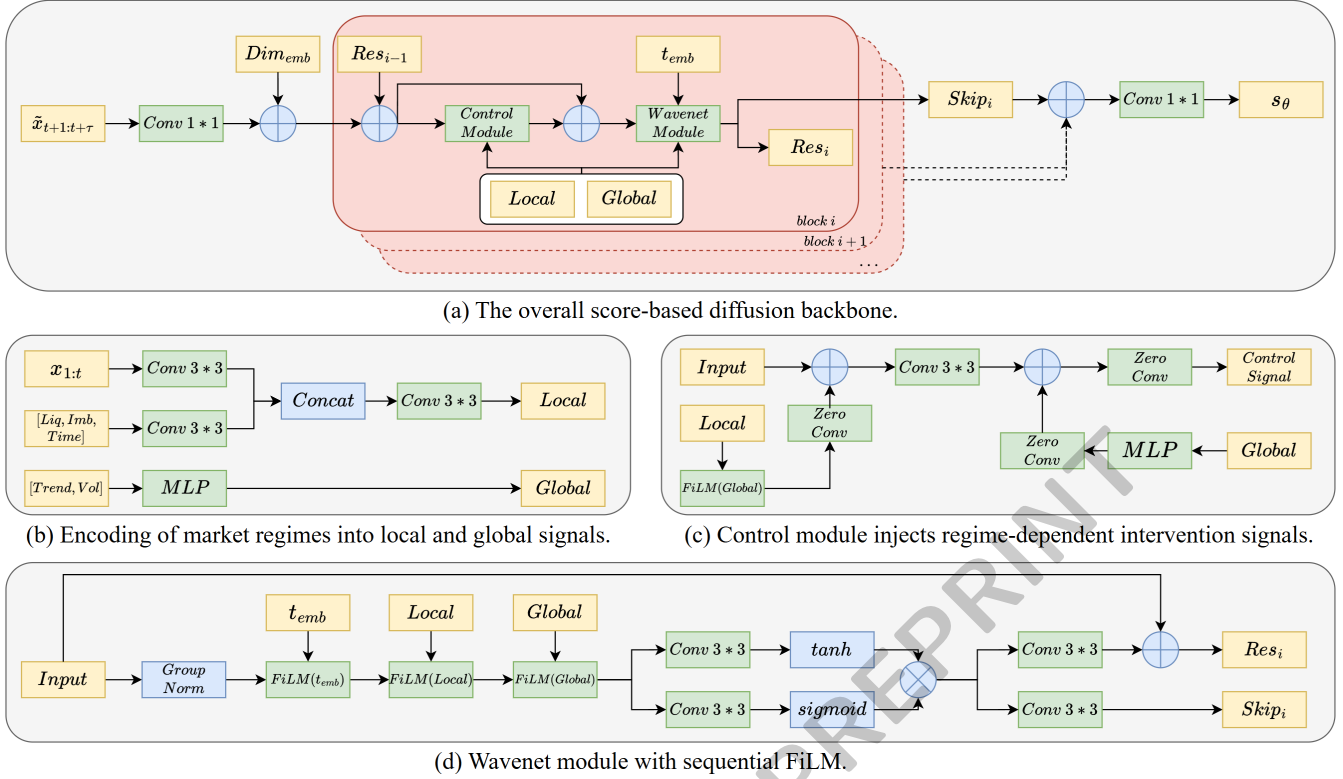


Figure 1: Illustration of DiffLOB Architecture.

represent longer-horizon market characteristics, such as future trend and volatility, and are encoded using multilayer perceptrons. This separation allows DiffLOB to model heterogeneous regime effects at different temporal scales.

Control Module and Dual-stage Training. To enforce controlling over generated trajectories, DiffLOB introduces a control module inspired by ControlNet [Zhang *et al.*, 2023], as shown in (Figure 1(c)). The condition encoder and wavenet module are first trained to model the base data distribution. Subsequently, the control module is trained while keeping the other backbone parameters frozen. Control signals are injected additively through zero-initialized 1×1 convolutional layers, whose weights and biases are initialized to zero. As a result, the control pathway has no effect at initialization and gradually learns regime-specific interventions in a stable and interpretable manner.

Wavenet Module with FiLM Modulation. Each residual block follows a Wavenet-style gated architecture with residual and skip connections, as shown in Figure 1(d). Conditioning information is injected using feature-wise linear modulation (FiLM) [Perez *et al.*, 2018]. Block activations are sequentially modulated by the diffusion timestep embedding t_{emb} , *Local* and *Global* regime embeddings. This sequential modulation scheme enables structured interaction between diffusion time, short-term market states, and long-horizon regime characteristics.

4 Experimental Setup

In this section, we provide data, preprocessing methods, as well as the training and sampling procedures.

4.1 Data and Preprocessing

We use the LOBSTER data² as our LOB data source [Huang and Polak, 2011]. We sample one snapshot per second. Each snapshot includes the top 10 levels on both the bid and ask sides. In particular, we train, evaluate, and test our model separately on three stocks, which are AMZN, AAPL and GOOG. For each stock, we use 16 consecutive trading days for training (1 February to 23 February 2023), 1 day for validation (24 February 2023), and the final 2 days for testing (27-28 February 2023).

Price Representation. Let $a_t^{(k)}$ and $b_t^{(k)}$ denote the ask and bid prices at level k at time t . We define the mid-price as

$$m_t = \frac{a_t^{(1)} + b_t^{(1)}}{2}.$$

To model temporal price dynamics, we use one-step mid-price difference:

$$r_t = m_{t+1} - m_t.$$

In addition, we include cross-sectional price differences to capture the spatial structure of the order book:

$$\Delta a_t^{(k)} = a_t^{(k)} - a_t^{(k-1)}, \quad \Delta s_t = a_t^{(1)} - b_t^{(1)}, \quad \Delta b_t^{(k)} = b_t^{(k-1)} - b_t^{(k)}$$

²<https://lobsterdata.com/>

for $k = 2, \dots, K$. The final price representation concatenates the one-step mid-price difference and the cross-sectional price differences. This construction preserves the original dimension while removing absolute price levels, resulting in a more stable and learnable representation for generative modeling.

Volume Representation. Raw volume v exhibit heavy-tailed distributions and large scale variations. To stabilize training, we cap volume values at the 99th percentile followed by a square-root normalization:

$$\tilde{v} = \frac{\sqrt{v}}{c},$$

where c is a constant scaling factor.

Market Regimes. Market regimes are computed from the target LOB trajectory $x_{t+1:t+\tau}$ and used as conditioning variables. Specifically, we define:

- **Trend** as the cumulative mid-price return over τ :

$$c_{t+1:t+\tau}^{\text{trend}} = \sum_{i=0}^{\tau-1} r(t+i).$$

- **Volatility** as the standard deviation of returns over τ :

$$c_{t+1:t+\tau}^{\text{vol}} = \sqrt{\frac{1}{\tau} \sum_{i=0}^{\tau-1} r_{t+i}^2 - \left(\frac{1}{\tau} \sum_{i=0}^{\tau-1} r_{t+i} \right)^2}.$$

- **Liquidity** as the total standing volume over each timestep:

$$c_{t+i}^{\text{liq}} = \sum_{k=1}^K \left(v_{t+i}^{a,(k)} + v_{t+i}^{b,(k)} \right).$$

- **Order-flow imbalance** as the normalized volume imbalance over each timestep:

$$c_{t+i}^{\text{imb}} = \frac{\sum_{k=1}^K \left(v_{t+i}^{a,(k)} - v_{t+i}^{b,(k)} \right)}{\sum_{k=1}^K \left(v_{t+i}^{a,(k)} + v_{t+i}^{b,(k)} \right)}.$$

4.2 Training and Sampling

Historical length t and generated length τ are fixed to 32 for all experiments. To promote stable training and reduce overfitting, we employ early stopping and exponential moving average (EMA) of model parameters. Early stopping is based on the validation loss and terminates training if no improvement greater than 0.001 is observed over 100 epochs. The model is optimized using Adam [Kingma and Ba, 2015] with a learning rate of 1×10^{-4} and a batch size of 128. We adopt the DDPM parameterization with 100 discrete noise levels and generate samples via ancestral sampling [Ho *et al.*, 2020]. Conditions are incorporated during training and sampling using classifier-free guidance [Ho and Salimans, 2021]. The network consists of 16 residual blocks with SiLU activations applied after each convolution.

5 Experimental Results

We evaluate DiffLOB from three complementary perspectives to assess its controllable realism, counterfactual validity, and practical usefulness. We compare DiffLOB with several baseline generative models. These include diffusion-based baselines (Diff-CSDI [Tashiro *et al.*, 2021] and Diff-S4 [Alcaraz and Strodthoff, 2022]), non-diffusion generative models (cGAN [Cont *et al.*, 2023] and cVAE [Sohn *et al.*, 2015]), and an autoregressive S4-based model AR [Nagy *et al.*, 2023]. DiffLOB without the control module is also included for ablation study use. For clarity of presentation, all figures use AMZN as an illustrative example.

5.1 Controllable Realism

We first evaluate the realism of DiffLOB by conditioning the model on *observed* future market regimes and comparing the generated trajectories with real data. The goal of this evaluation is to assess whether DiffLOB can faithfully reproduce key statistics of LOB dynamics.




Figure 2: Realism on Price.

Figure 2 evaluates the realism of price dynamics generated by models. Figure 2a shows that DiffLOB closely matches the empirical distributions of mid-price returns across multiple horizons. Figure 2b reports the distribution of bid-ask spread, where DiffLOB replicates real distribution best. Figure 2c compares the autocorrelation of absolute returns. DiffLOB reproduces the persistent decay pattern characteristic of volatility clustering, whereas cGAN and cVAE largely fail to capture temporal dependence and AR models underestimate long-range correlations.

Figure 3 is calculated from the first-order differences of the volume snapshots ($v_{t+1} - v_t$). A notable characteristic in the temporal difference correlation structure is the negative correlation at adjacent price levels, which is captured only by DiffLOB.

	AMZN				AAPL				GOOG			
Price-Realism	KS	Wasserstein	KL	JS	KS	Wasserstein	KL	JS	KS	Wasserstein	KL	JS
DiffLOB	0.052384	0.028371	0.075605	0.017585	0.031695	0.016446	0.026561	0.005704	0.055462	0.035195	0.363414	0.037249
DiffLOB w/o C	0.212398	0.182775	0.730008	0.064836	0.079402	0.064358	0.095074	0.014745	0.103601	0.085803	0.30239	0.027799
Diff-CSDI	0.628793	64.621056	5.862505	0.306881	0.576959	68.690042	5.507347	0.286949	0.573082	47.207406	4.99384	0.261333
Diff-S4	0.200429	0.155586	0.304272	0.051942	0.101795	0.0915	0.149292	0.019214	0.153828	0.06162	0.603084	0.058418
CGAN	0.053521	0.048909	0.378863	0.026156	0.830413	2.402774	9.343481	0.466478	0.28967	0.296793	1.523045	0.114271
CVAE	0.173396	0.142745	0.406215	0.053637	0.025862	0.021286	0.135777	0.013759	0.155807	0.085175	0.419525	0.050767
AR	0.118031	0.076163	0.419127	0.037131	0.173451	0.145696	0.304234	0.03807	0.107334	0.046178	0.324453	0.040902
Volume-Realism	KS	Wasserstein	KL	JS	KS	Wasserstein	KL	JS	KS	Wasserstein	KL	JS
DiffLOB	0.109174	97.54283	0.111954	0.023297	0.087666	55.377778	0.071645	0.014677	0.099899	97.270339	0.086374	0.016609
DiffLOB w/o C	0.113053	117.36868	0.108419	0.022918	0.100573	72.929232	0.091735	0.019714	0.126238	132.065501	0.117649	0.023473
Diff-CSDI	0.293449	599243.7598	0.324422	0.052769	0.235838	587451.9702	0.29902	0.04794	0.256603	578472.3812	0.266948	0.039498
Diff-S4	0.170321	161.665895	0.178012	0.036454	0.132301	104.168332	0.138386	0.025026	0.246833	279.02256	0.325909	0.060015
CGAN	0.288535	321.049535	0.500225	0.090229	0.487548	275.698425	2.956931	0.23594	0.127321	151.897776	0.137742	0.027589
CVAE	0.260795	255.043006	1.154918	0.153781	0.23728	169.905872	1.251494	0.140914	0.231203	240.882982	0.864015	0.112463
AR	0.397229	444.216686	1.172602	0.159319	0.285415	188.199251	0.669676	0.112751	0.2863	343.938149	0.724715	0.106982

Table 1: Controllable Realism on Three Stocks.




Figure 3: Temporal Difference Volume Correlation.

Figure 4 visualizes the distributions of future market regime variables under observed conditions, including trend, volatility, liquidity, and order-flow imbalance. DiffLOB accurately captures the central tendency of trend and volatility, while achieving an excellent match to the full empirical distributions of liquidity and order-flow imbalance, including both spread and tail behavior. In contrast, baseline models exhibit noticeable mismatches, such as overly concentrated distributions (AR on Trend) and shifted modes (cVAE on Liquidity, AR on Volatility, AR on Imbalance). These results further confirm that DiffLOB can faithfully reproduce regime-level statistics when conditioned on true future regimes, supporting the quantitative findings in Table 1.

Table 1 reports quantitative distributional distances between generated and real LOB trajectories under observed future regimes on the three stocks across both price-related and volume-related statistics. We consider four distance metrics, including Kolmogorov–Smirnov (KS) statistic, Wasserstein distance, Kullback–Leibler (KL) divergence, and Jensen–Shannon (JS) divergence. Across all assets and metrics, DiffLOB consistently achieves the lowest or near-lowest distances, indicating that it best matches the empirical distributions of real data when conditioned on true future regimes. Removing the control module (DiffLOB w/o C) leads to a clear degradation in performance, highlighting the importance of explicit regime control even under observed conditions. Diffusion-based baselines (Diff-CSDI and Diff-S4) improve upon non-diffusion methods in some cases, but remain substantially less accurate than DiffLOB, show-




Figure 4: Controllable Realism Distribution..

ing the effectiveness of our proposed DiffLOB architecture described in Section 3.2. Non-diffusion baselines, including cGAN, cVAE, and the autoregressive model, exhibit significantly larger discrepancies, especially in Wasserstein and KL distances. Overall, these results demonstrate that DiffLOB achieves superior controllable realism across both price and volume dimensions, and that designed control module plays a critical role in accurately reproducing LOB statistics.

5.2 Counterfactual Validity

We next evaluate the counterfactual validity by intervening on future market regimes and examining the resulting generated trajectories. Specifically, we impose extreme hypothetical conditions, including high and low trend, volatility, liquidity, and order-flow imbalance, and assess whether the generated samples exhibit statistical properties consistent with the imposed regimes.




Figure 5: Counterfactual Realism Distribution.

Figure 5 evaluates counterfactual validity under explicit regime interventions. For each regime variable, we compare distributions generated under high and low counterfactual conditions with real trajectories in the corresponding extreme regimes (top and bottom 20%). For trend and volatility, DiffLOB does not perfectly overlap with empirical distributions but consistently shifts in the correct direction and preserves the separation between high and low conditions. In contrast, for liquidity and order-flow imbalance, DiffLOB closely matches the empirical distributions, indicating strong controllability for volume-driven dynamics. Overall, these results show that DiffLOB learns an interpretable and intervention-consistent relationship between future regimes and LOB dynamics, enabling meaningful counterfactual generation.

Figure 6 and Figure 7 provide visualizations of counterfactual generation on price and volume trajectories. In Figure 6, explicit interventions on future trend and volatility induce coherent and interpretable changes. High-trend conditions lead to persistent upward price movements, while low-trend conditions result in declining trajectories. Simi-




Figure 6: Counterfactual LOB Price Trajectories.

larly, high-volatility interventions produce more volatile price paths with larger fluctuations, whereas low-volatility conditions yield smoother dynamics. In Figure 7, imposing high-




Figure 7: Counterfactual LOB Volume Trajectories.

liquidity conditions leads to consistently larger volume magnitudes across the LOB, while low-liquidity conditions shift the overall volume downward. Imbalance interventions induce clear asymmetries between bid and ask sides, consistent with the imposed order-flow imbalance. These results demonstrate that DiffLOB can generate regime-consistent dynamics beyond matching marginal distributions.

Table 2 reports quantitative distances computed between counterfactual samples generated with a fixed future regime (e.g., high trend) and real market trajectories whose regime values fall into the same extreme quantile (top or bottom 20%). cGAN performs competitively on trend-related metrics but degrades substantially for volatility, liquidity, and imbalance, indicating limited regime robustness. DiffLOB without the control module performs relatively well under high-liquidity and high-imbalance conditions, suggesting that

High-Trend	KS	Wasserstein	KL	JS
DiffLOB	0.444095	0.522512	2.517532	0.180279
DiffLOB w/o C	0.515225	0.801835	4.879508	0.265181
Diff-CSDI	0.627345	65.274135	5.855148	0.307245
Diff-S4	0.292354	0.264849	1.098681	0.092509
CGAN	0.048718	0.036091	0.422542	0.026361
CVAE	0.298339	0.225797	0.79771	0.092942
AR	0.200438	0.199092	1.210306	0.091039
Low-Trend	KS	Wasserstein	KL	JS
DiffLOB	0.425788	0.507738	3.529521	0.182637
DiffLOB w/o C	0.339051	0.36766	2.525023	0.141955
Diff-CSDI	0.755996	84.376958	6.563037	0.356456
Diff-S4	0.431509	0.521119	3.855445	0.187511
CGAN	0.039677	0.044298	0.47366	0.023752
CVAE	0.38066	0.459526	2.89087	0.156796
AR	0.264977	0.296113	1.946188	0.107689
High-Volatility	KS	Wasserstein	KL	JS
DiffLOB	0.083474	0.06075	0.169118	0.03113
DiffLOB w/o C	0.292613	0.237796	1.004272	0.108261
Diff-CSDI	0.609586	59.726185	5.559404	0.286991
Diff-S4	0.118421	0.12436	0.67194	0.064088
CGAN	0.178244	0.088973	0.524627	0.048186
CVAE	0.118761	0.077975	0.646276	0.071165
AR	0.083668	0.065337	0.454682	0.041336
Low-Volatility	KS	Wasserstein	KL	JS
DiffLOB	0.150047	0.072569	0.772114	0.061085
DiffLOB w/o C	0.163749	0.140627	1.38697	0.075512
Diff-CSDI	0.636904	59.84074	5.663091	0.293081
Diff-S4	0.297515	0.160111	0.999069	0.111148
CGAN	0.16065	0.113201	1.285274	0.078348
CVAE	0.287089	0.16455	1.092421	0.100044
AR	0.244042	0.141	1.457307	0.09852
High-Liquidity	KS	Wasserstein	KL	JS
DiffLOB	0.137538	168.024319	0.208135	0.034988
DiffLOB w/o C	0.134555	179.447363	0.18559	0.03162
Diff-CSDI	0.318472	617445.7056	0.419245	0.05283
Diff-S4	0.199552	229.470802	0.279902	0.050367
CGAN	0.389481	513.984844	1.1278	0.153755
CVAE	0.335492	365.126803	1.982426	0.229481
AR	0.422048	445.418425	1.840147	0.215084
Low-Liquidity	KS	Wasserstein	KL	JS
DiffLOB	0.126894	104.19346	0.099957	0.021029
DiffLOB w/o C	0.140489	110.790344	0.114877	0.025278
Diff-CSDI	0.311335	585813.0602	0.444758	0.073959
Diff-S4	0.224048	142.265977	0.212572	0.048326
CGAN	0.214546	167.358666	0.315572	0.063397
CVAE	0.351886	231.117557	1.748266	0.220437
AR	0.410027	274.375898	1.572332	0.19979
High-Imbalance	KS	Wasserstein	KL	JS
DiffLOB	0.142787	154.156402	0.182522	0.032153
DiffLOB w/o C	0.130096	155.84261	0.181721	0.032042
Diff-CSDI	0.303622	629928.7738	0.363013	0.054523
Diff-S4	0.163933	165.329592	0.180472	0.035623
CGAN	0.285831	342.915205	0.59001	0.095159
CVAE	0.274923	291.958093	1.485293	0.171144
AR	0.432061	397.430331	1.656605	0.202636
Low-Imbalance	KS	Wasserstein	KL	JS
DiffLOB	0.138717	147.853412	0.181154	0.034248
DiffLOB w/o C	0.153997	174.322966	0.204804	0.038629
Diff-CSDI	0.307549	593616.0608	0.43386	0.063705
Diff-S4	0.198524	209.414158	0.253176	0.049452
CGAN	0.298427	356.49669	0.635989	0.101998
CVAE	0.290173	290.269725	1.529841	0.18089
AR	0.383488	338.46483	1.580578	0.186188

Table 2: Counterfactual Validity on AMZN.

volume-driven regimes can be partially captured by the diffusion backbone alone, consistent with findings in [Wang and Ventre, 2025]. However, across the majority of regimes and distance metrics, DiffLOB consistently achieves the lowest or near-lowest distances. This demonstrates that explicitly modeling future regimes through the control module leads to more reliable and coherent counterfactual responses across diverse market conditions.

5.3 Counterfactual Usefulness

Finally, we assess the practical usefulness of counterfactual trajectories generated by DiffLOB. We consider two tasks: trend prediction and liquidity prediction. Trend prediction is formulated as a classification task that predicts the direction of future price movement, and is evaluated using accuracy (Acc). Liquidity prediction is treated as a regression task that predicts future liquidity values, evaluated using the coefficient of determination (R^2). For both tasks, models are trained on past 1 minute LOB data to predict future 1 minute regime. Models are evaluated on 1 March 2023, with performance reported separately on top 20% and bottom 20% of the regime distribution. We compare three training settings: *Real*, using only real samples; *Real * 2*, duplicating real data to control for dataset size; and *Real+CF*, augmenting real data with counterfactual trajectories to enrich extreme regime coverage.

	AMZN		AAPL		GOOG	
Trend Prediction	Acc-High	Acc-Low	Acc-High	Acc-Low	Acc-High	Acc-Low
Real	0.983	0.061	0.885	0.14	0.517	0.453
Real * 2	0.98	0.064	0.89	0.153	0.5	0.488
Real + CF	0.944	0.289	0.918	0.17	0.56	0.556
Liquidity Prediction	R ² -High	R ² -Low	R ² -High	R ² -Low	R ² -High	R ² -Low
Real	-1.568	0.086	-2.512	-0.184	-0.742	-0.039
Real * 2	-1.656	0.127	-2.525	-0.26	-0.796	-0.011
Real + CF	-1.509	0.278	-1.432	0.17	0.017	-0.053

Table 3: Counterfactual Usefulness on the three Stocks.

In Table 3, we can see that augmenting real data with counterfactual samples improves performance on most high- and low-regime subsets, compared to training on real data alone. The gains demonstrate that counterfactual trajectories provide complementary information that is scarce in real data, proving the usefulness of counterfactual trajectories by DiffLOB.

6 Conclusion

We introduce DiffLOB, a diffusion-based framework for controllable and counterfactual generation of limit order book trajectories. By explicitly conditioning on future market regimes—including trend, volatility, liquidity, and order-flow imbalance—DiffLOB enables direct intervention on hypothetical future conditions. Extensive experiments demonstrate that DiffLOB achieves superior controllable realism, generates coherent and regime-consistent counterfactual trajectories, and provides tangible benefits for downstream prediction tasks under extreme market regimes. These results highlight the importance of explicit regime-aware control for realistic simulation and meaningful counterfactual analysis.

Acknowledgments

Carmine Ventre was partially supported by the Engineering and Physical Sciences Research Council (EPSRC) (Grant numbers EP/W034042/1 and UKRI4151).

References

- [Alcaraz and Strodthoff, 2022] Juan Miguel Lopez Alcaraz and Nils Strodthoff. Diffusion-based time series imputation and forecasting with structured state space models. *Transactions on Machine Learning Research*, 2022.
- [Berti *et al.*, 2025] Leonardo Berti, Bardh Prenkaj, and Paola Velardi. Trades: Generating realistic market simulations with diffusion models. *Proceedings of the Twenty-Eighth European Conference on Artificial Intelligence*, 2025.
- [Briola *et al.*, 2025a] Antonio Briola, Silvia Bartolucci, and Tomaso Aste. Deep limit order book forecasting: a microstructural guide. *Quantitative Finance*, 25(7):1101–1131, 2025.
- [Briola *et al.*, 2025b] Antonio Briola, Silvia Bartolucci, and Tomaso Aste. Hlob–information persistence and structure in limit order books. *Expert Systems with Applications*, 266:126078, 2025.
- [Coletta *et al.*, 2021] Andrea Coletta, Matteo Prata, Michele Conti, Emanuele Mercanti, Novella Bartolini, Aymeric Moulin, Svitlana Vyetrenko, and Tucker Balch. Towards realistic market simulations: a generative adversarial networks approach. In *Proceedings of the Second ACM International Conference on AI in Finance*, pages 1–9, 2021.
- [Cont *et al.*, 2023] Rama Cont, Mihai Cucuringu, Jonathan Kochems, and Felix Prenzel. Limit order book simulation with generative adversarial networks. *Available at SSRN 4512356*, 2023.
- [Gould *et al.*, 2013] Martin D Gould, Mason A Porter, Stacy Williams, Mark McDonald, Daniel J Fenn, and Sam D Howison. Limit order books. *Quantitative Finance*, 13(11):1709–1742, 2013.
- [Hashimoto *et al.*, 2025] Ryuji Hashimoto, Yuki Tanaka, Takehiro Takayanagi, Zhe Piao, and Kiyoshi Izumi. Norm-salvaged embedding: Improving condition alignment of synthetic time series generation in finance. In *Proceedings of the 6th ACM International Conference on AI in Finance*, pages 838–846, 2025.
- [Ho and Salimans, 2021] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- [Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [Huang and Polak, 2011] Ruihong Huang and Tomas Polak. Lobster: Limit order book reconstruction system. *Available at SSRN 1977207*, 2011.
- [Hultin *et al.*, 2023] Hanna Hultin, Henrik Hult, Alexandre Proutiere, Samuel Samama, and Ala Tarighati. A generative model of a limit order book using recurrent neural networks. *Quantitative Finance*, 23(6):931–958, 2023.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015.
- [Koa *et al.*, 2023] Kelvin JL Koa, Yunshan Ma, Ritchie Ng, and Tat-Seng Chua. Diffusion variational autoencoder for tackling stochasticity in multi-step regression stock price prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1087–1096, 2023.
- [Li *et al.*, 2025] Junjie Li, Yang Liu, Weiqing Liu, Shikai Fang, Lewen Wang, Chang Xu, and Jiang Bian. Mars: a financial market simulation engine powered by generative foundation model. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [Nagy *et al.*, 2023] Peer Nagy, Sascha Frey, Silvia Saporá, Kang Li, Anisoara Calinescu, Stefan Zohren, and Jakob Foerster. Generative ai for end-to-end limit order book modelling: A token-level autoregressive generative model of message flow using a deep state space network. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, pages 91–99, 2023.
- [Perez *et al.*, 2018] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [Sohl-Dickstein *et al.*, 2015] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [Sohn *et al.*, 2015] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015.
- [Song and Ermon, 2019] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [Song *et al.*, 2021] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *International Conference on Learning Representations*, 2021.
- [Tanaka *et al.*, 2025] Yuki Tanaka, Ryuji Hashimoto, Takehiro Takayanagi, Zhe Piao, Yuri Murayama, and Kiyoshi Izumi. Cofindiff: Controllable financial diffusion model for time series generation. *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence*, 2025.

- [Tashiro *et al.*, 2021] Yusuke Tashiro, Jiaming Song, Yang Song, and Stefano Ermon. Csd: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems*, 34:24804–24816, 2021.
- [van den Oord *et al.*, 2016] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew W. Senior, and Koray Kavukcuoglu. WaveNet: A Generative Model for Raw Audio. In *Proceedings of the 9th ISCA Speech Synthesis Workshop (SSW)*, pages 125–131, 2016.
- [Vincent, 2011] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [Wang and Ventre, 2024] Zhuohan Wang and Carmine Ventre. A financial time series denoiser based on diffusion models. In *Proceedings of the 5th ACM International Conference on AI in Finance*, pages 72–80, 2024.
- [Wang and Ventre, 2025] Zhuohan Wang and Carmine Ventre. Diffvolume: Diffusion models for volume generation in limit order books. In *Proceedings of the 6th ACM International Conference on AI in Finance*, pages 587–595, 2025.
- [Zhang *et al.*, 2023] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847, 2023.