

Controllable Financial Market Generation with Diffusion Guided Meta Agent

Yu-Hao Huang¹, Chang Xu^{2*}, Yang Liu², Weiqing Liu², Wu-Jun Li¹, Jiang Bian²

¹National Key Laboratory for Novel Software Technology, School of Computer Science, Nanjing University

²Microsoft Research Asia

huangyh@smail.nju.edu.cn, {chanx, yangliu2, weiqing.liu}@microsoft.com,
liwujun@nju.edu.cn, jiang.bian@microsoft.com

Abstract

Generative modeling has transformed many fields, such as language and visual modeling, while its application in financial markets remains under-explored. As the minimal unit within a financial market is an order, order-flow modeling represents a fundamental generative financial task. However, current approaches often yield unsatisfactory fidelity in generating order flow, and their generation lacks controllability, thereby limiting their practical applications. In this paper, we formulate the challenge of controllable financial market generation, and propose a *Diffusion Guided Meta Agent* (DigMA) model to address it. Specifically, we employ a conditional diffusion model to capture the dynamics of the market state represented by time-evolving distribution parameters of the mid-price return rate and the order arrival rate, and we define a meta agent with financial economic priors to generate orders from the corresponding distributions. Extensive experimental results show that DigMA achieves superior controllability and generation fidelity. Moreover, we validate its effectiveness as a generative environment for downstream high-frequency trading tasks and its computational efficiency.

Extended version — <https://arxiv.org/abs/2408.12991>

Introduction

Generative modeling has transformed fields such as natural language processing (Brown et al. 2020; Chowdhery et al. 2023; Singhal et al. 2023), media synthesis (Romach et al. 2022; Lu et al. 2024; Copet et al. 2024), science discovery (Wang et al. 2021; Skinnider et al. 2021; M. Bran et al. 2024; Fürtter, Muñoz-Gil, and Briegel 2024) and medical applications (Aversa et al. 2023; Tu et al. 2024), with techniques such as Generative Adversarial Networks (GANs) (Goodfellow et al. 2014), Diffusion Models (DMs) (Ho, Jain, and Abbeel 2020), and pre-trained Transformers (Vaswani 2017) serving as core building blocks. While the financial market is a data-intensive and intricate system characterized by dynamic interactions among participants and has attracted research attention for decades, the application of generative models into this domain remains rare. Similar to words in language and pixels

in images, an order is the fundamental element representing the minimal unit of events within a financial market (Palmer et al. 1994; Lux and Marchesi 1999; Raberto et al. 2001; Chiarella and Iori 2002; Chiarella, Iori, and Perello 2009). Through sequential order flow, researchers, investors, and policymakers can investigate the intrinsic interactive mechanisms and microstructure of financial markets (Raberto et al. 2001; Mizuta 2020; Coletta et al. 2021; Guo et al. 2023). Therefore, order-flow modeling constitutes a fundamental generative task in financial domains.

Recent works have attempted to simulate order-level financial markets using agent-based methods, employing either rule-based agents (Vyetenko et al. 2020; Byrd, Hybinette, and Balch 2020; Amrouni et al. 2021) or learned agents (Li et al. 2020; Coletta et al. 2023; Li et al. 2025b). However, their fidelity and flexibility remain limited. Rule-based agents rely on over-simplified assumptions about the market and can only represent known types of market participants under predefined scenario. They are not trained on real market data but are instead constructed using hand-crafted rules, resulting in limited simulation fidelity. Learned agents are trained to predict next order given history order flow, where the order flow may contain hundreds of orders within a single minute. It is challenging to directly capture the long term dependency, as a trading day spans hundreds of minutes. These models tend to emphasize the local distribution of the simulated order flow while neglecting the global dynamic. More importantly, controllability of the generated market, which would enable researchers and practitioners to systematically explore market behaviors under various conditions such as extreme or rare events, is absent in the literature. This underscores a practical gap for conducting scenario-based experiments and counterfactual analysis (Mizuta 2020; Guo et al. 2023).

In contrast to existing works, we propose incorporating controllability into financial market modeling by formulating the problem as a conditional generation task. The objective is to construct specific scenarios characterized by varying levels of asset return, intraday volatility, and rare events such as sharp drops or extreme amplitudes. Achieving this requires addressing the critical challenge of establishing a principled connection between high-level control targets, such as desired market scenarios, and the orders generated by the model. Moreover, exercising control at the order level

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

is particularly challenging, because (1) the long and irregular length of real-world order-flow sequences hinders the direct application of diffusion models to raw order-level data, and (2) linking a “macro” control target to each individual “micro” order is impractical due to the inherently low signal-to-noise ratio present in order flow.

In this paper, we formulate the challenge of controllable financial market generation, and propose the *Diffusion Guided Meta Agent* (DigMA) model to address the challenge. Specifically, we employ a conditional diffusion model to capture the dynamics of the market state represented by time-evolving distribution parameters of the mid-price return rate and the order arrival rate, and we define a meta agent with financial economic priors to generate orders from the corresponding distributions. With DigMA, we are able to control the generation process to simulate order flow under target scenario with high fidelity. Our contributions can be summarized as follows:

- We formulate controllable financial market generation problem as a novel and practically meaningful challenge for machine learning application in finance.
- To the best of our knowledge, DigMA is among the pioneering models to integrate advanced diffusion-based generative techniques into financial market modeling.
- Through extensive experiments on real stock market data, we demonstrate that DigMA achieves superior controllability and the highest fidelity to established stylized facts. Moreover, we validate its effectiveness as a generative environment for downstream high-frequency trading tasks, as well as its computational efficiency.

Preliminaries and Related Work

In this section, we briefly introduce preliminaries of the limit order book and review related research on financial market simulation and diffusion models.

Limit Order Book

The majority of financial markets worldwide operate under a double-auction mechanism, in which orders serve as the fundamental units of events. An order consists of four basic elements: timestamp t , price p , quantity q , and order type o . Real markets feature a variety of order types, such as limit orders, market orders, cancel orders, and conditional orders. In the literature, it is often sufficient to represent trading decisions using limit orders and cancel orders (Chiarella, Iori, and Perello 2009).

The output of a market simulation model is a sequence of orders $\mathcal{O} = \{(t_1, p_1, q_1, o_1), (t_2, p_2, q_2, o_2), \dots\}$, commonly referred to as the order flow, which collectively determines the order book and the resulting price series. The order book is defined as the set of outstanding orders that have not yet been executed. Limit orders can be further categorized into buy and sell limit orders, whose prices are referred to as bids and asks, respectively. An order book consisting exclusively of limit orders is known as the limit order book (LOB). The transaction price at each timestamp is typically defined as the price of the most recently executed order. Sampling these prices at a chosen frequency yields the

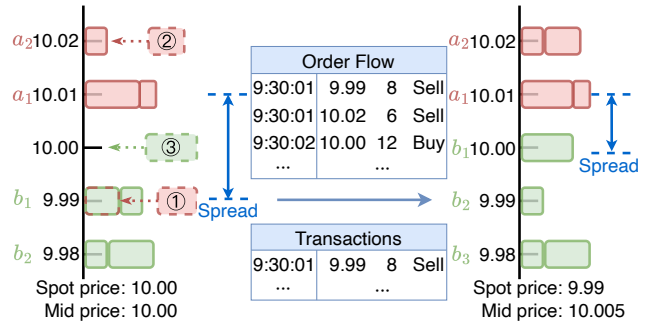


Figure 1: Limit order book and order flow

corresponding price series. An illustrative example of a limit order book is shown in Figure 1.

Financial Market Simulation

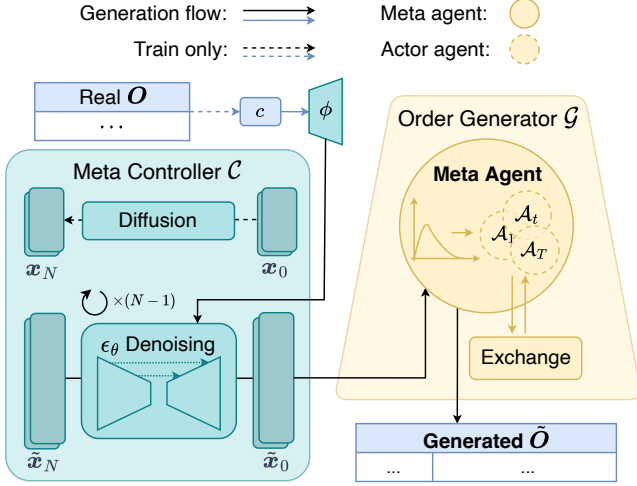
Early work on financial market simulation adopted a multi-agent approach (Raberto et al. 2001; Palmer et al. 1994; Lux and Marchesi 1999), using rule-based agents under simplified trading protocols to replicate “stylized facts” (Cont 2001), such as volatility clustering. Chiarella and Iori (2002); Chiarella, Iori, and Perello (2009) extended these simulations to order-driven markets, aligning more closely with modern stock market structures. Subsequent studies further customized agent behaviors within this framework to better support research on decision-making (Vyetenko et al. 2020; Byrd, Hybinette, and Balch 2020; Wang and Wellman 2017). With recent advances in machine learning, researchers have also explored neural networks as world agents capable of directly predicting orders from historical order flow (Coletta et al. 2021; Li et al. 2020; Coletta et al. 2022, 2023). In contrast, our model employs a conditional diffusion model to guide an agent-based system in order generation.

Diffusion Models

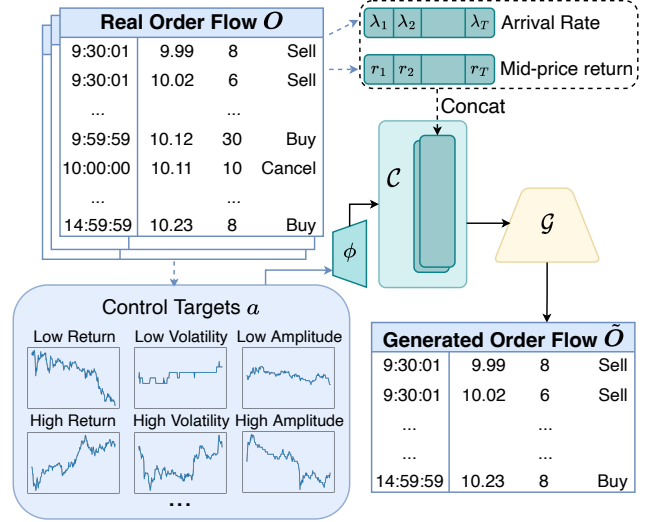
Diffusion probabilistic models, commonly referred to as diffusion models, learn to transform between simple and target distributions through a sequence of small perturbations defined by a diffusion process (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020; Song et al. 2021). These models have been applied to data generation across a wide range of modalities, including images (Rombach et al. 2022; Aversa et al. 2023; Dhariwal and Nichol 2021; Song et al. 2023), audio (Kong et al. 2021; Liu et al. 2023), video (Brooks et al. 2024; Guo et al. 2024), and general time series (Kolloviev et al. 2023; Fan et al. 2024; Yuan and Qiao 2024; Huang et al. 2025; Li et al. 2025a; Deng et al. 2025). In contrast to existing work, we are the first to apply diffusion models to the generation of financial orders.

Method

In this section, we first present the problem formulation for controllable financial market generation. We then describe the architecture of our *Diffusion Guided Meta Agent* (DigMA) model in detail.



(a) Meta controller and order generator in DigMA.



(b) Data flow of DigMA.

Figure 2: Overview of DigMA model. Raw order flow is processed into market states for the meta controller to learn. The meta agent is guided by the meta controller, and generates simulated order flow.

Problem Formulation: Controllable Financial Market Generation

Unlike standard market simulation, which generates order flow unconditionally, controllable market generation aims to simulate order flows corresponding to specified scenarios using generative models. Such scenarios can be characterized by indicators representing aggregated order-flow statistics, including daily return, daily amplitude, and intraday volatility. Formally, let \mathcal{F} denote a function that computes a given indicator. For a real order-flow sample $\mathcal{O} \sim q(\mathcal{O})$, the corresponding indicator value a is computed as $a = \mathcal{F}(\mathcal{O})$.

In controllable financial market generation, a market generator \mathcal{M} defines a conditional sampler $p_{\mathcal{M}}(\mathcal{O} | a)$. Given a control target a , the generator produces an order-flow sample $\tilde{\mathcal{O}} \sim p_{\mathcal{M}}(\mathcal{O} | a)$. The controllability objective is to minimize the discrepancy between the target a and the indicator value of the generated order flow, $\tilde{a} = \mathcal{F}(\tilde{\mathcal{O}})$. Formally, the controllability objective is given by

$$\min_{\mathcal{M}} \mathbb{E}_{a, \tilde{a}} [\|\tilde{a} - a\|^2] = \min_{\mathcal{M}} \mathbb{E}_{a, \tilde{\mathcal{O}}} [\|\mathcal{F}(\tilde{\mathcal{O}}) - a\|^2]. \quad (1)$$

In addition, the generator is required to produce order flows with high fidelity. This objective aligns with the general goal of market simulation, which is to minimize a divergence measure \mathcal{D} between the distributions of “stylized facts” computed from real and generated order flows. Here, the stylized facts can also be expressed as a set of aggregated statistics with corresponding computation functions. Let \mathcal{F}' denote a function that computes a stylized fact, and let $p(\cdot)$ denote a probability density. The fidelity objective is then defined as

$$\min_{\mathcal{M}} \mathcal{D} \left(p(\mathcal{F}'(\tilde{\mathcal{O}})) \parallel p(\mathcal{F}'(\mathcal{O})) \right). \quad (2)$$

Diffusion Guided Meta Agent Model

Order flow data are highly intricate and noisy, often consisting of tens of thousands of events per stock daily, which

poses significant computational challenges. Modeling the full distribution of order flow across diverse scenarios is therefore non-trivial. Moreover, directly linking a “macro” control target to individual “micro” orders is impractical due to the low signal-to-noise ratio of order flow. To address these challenges, we adopt a two-stage design instead of applying diffusion models directly to raw order flow.

DigMA consists of two modules. The first is a meta controller \mathcal{C} , which learns the intraday dynamics of latent market states \mathbf{x} conditioned on a scenario c , and models the distribution $q(\mathbf{x} | c)$ using a conditional diffusion model. The second is an order generator \mathcal{G} , which comprises a simulated exchange and a meta agent. The meta agent incorporates financial economic priors and is guided by the meta controller to generate orders through a stochastic process. An overview of the DigMA architecture is shown in Figure 2. Overall, the model is expressed as $\mathcal{M} = \{\mathcal{C}, \mathcal{G}\}$.

Meta Controller To represent intraday market dynamics associated with a given scenario c , the market states \mathbf{x} are required to evolve over time while maintaining a causal relationship with c . We therefore extract the minutely mid-price return rate \mathbf{r} and the order arrival rate $\boldsymbol{\lambda}$ from real order-flow data, and define the market states as $\mathbf{x} = \{\mathbf{r}, \boldsymbol{\lambda}\}$.

Since the number of trading minutes within each trading day is fixed, the stacked market-state sequence corresponding to a single day naturally forms one training sample. The objective is to fit the distribution of these market states:

$$\min \mathbb{E}_{\mathbf{x}} \mathcal{D}(p_{\mathcal{C}}(\mathbf{x}) \parallel q(\mathbf{x})). \quad (3)$$

Given training samples $\{\mathbf{x} \sim q(\mathbf{x})\}$, we generate noisy latent variables as $\mathbf{x}_n = \sqrt{\bar{\alpha}_n} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_n} \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and n denotes the diffusion step. Here, $\bar{\alpha}_n$ is derived from the variance schedule $\{\beta_n \in (0, 1)\}_{n=1}^N$, where $\alpha_n = 1 - \beta_n$ and $\bar{\alpha}_n = \prod_{i=1}^n \alpha_i$, and N denotes the maximum diffusion step. We adopt the ϵ -parameterized denoising diffusion probabilistic model (DDPM) (Ho, Jain, and

Abbeel 2020) and train it to predict the injected noise from \mathbf{x}_n as $\hat{\epsilon} = \epsilon_\theta(\mathbf{x}_n, n)$, where θ denotes the model parameters. The resulting training loss is defined as

$$L_M := \mathbb{E}_{\mathbf{x}, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), n} [\|\epsilon - \epsilon_\theta(\mathbf{x}_n, n)\|^2]. \quad (4)$$

During sampling, the diffusion process proceeds iteratively from \mathbf{x}_N to $\tilde{\mathbf{x}}_0$ following

$$\tilde{\mathbf{x}}_{n-1} = \frac{1}{\sqrt{\alpha_n}} \left(\tilde{\mathbf{x}}_n - \frac{1 - \alpha_n}{\sqrt{1 - \alpha_n}} \epsilon_\theta(\tilde{\mathbf{x}}_n, n) \right) + \sigma_n \mathbf{z}, \quad (5)$$

where $\mathbf{x}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and $\sigma_n = \sqrt{\beta_n}$. The resulting $\tilde{\mathbf{x}}_0$ provides a latent market trajectory that enables the meta controller to guide the meta agent during the subsequent order generation process.

To further enable controlled order generation, we condition the diffusion model on target scenario variables so as to steer the generated order flow toward the desired market regime. Following common practice (Rombach et al. 2022), we implement a conditional ϵ -parameterized noise predictor $\epsilon_\theta(\mathbf{x}_n, n, c)$ to support sampling under a specified control target c . Specifically, we adopt three indicators that are commonly used to characterize financial market states as control targets: daily return, amplitude, and volatility. Each of these indicators can be computed from the return rates of the price series, and controlling them allows the generated order flow to reflect a wide range of realistic market scenarios.

To incorporate control targets into the diffusion model, we introduce a target-specific feature extractor ϕ that projects target indicators into latent representations $\phi(c)$. We design two types of control encoders. The first is a *discrete* control encoder, where target conditions are mapped into predefined bins that are treated as class labels and embedded via a learnable embedding matrix. The second is a *continuous* control encoder, which uses a fully connected network to map real-valued conditions into latent vectors. The feature extractor ϕ is trained jointly with the conditional sampler using

$$L_C := \mathbb{E}_{\mathbf{x}, c, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), n} [\|\epsilon - \epsilon_\theta(\mathbf{x}_n, n, \phi(c))\|^2]. \quad (6)$$

For both encoders, we employ classifier-free guidance (Ho and Salimans 2022) to perform control. During training, unconditional and conditional samplers are jointly optimized by randomly dropping the conditioning information. During sampling, the guided score is computed as

$$\tilde{\epsilon}_{\theta, \phi}(\mathbf{x}_n, n, c) = (1 - s)\epsilon_\theta(\mathbf{x}_n, n) + s\epsilon_\theta(\mathbf{x}_n, n, \phi(c)), \quad (7)$$

where s is a scaling factor that controls the guidance strength. Sampling then follows Equation 5 with

$$\mathbf{x}_{n-1} = \frac{1}{\sqrt{\alpha_n}} \left(\mathbf{x}_n - \frac{1 - \alpha_n}{\sqrt{1 - \alpha_n}} \tilde{\epsilon}_{\theta, \phi}(\mathbf{x}_n, n, c) \right) + \sigma_n \mathbf{z}. \quad (8)$$

In practice, we adopt DDIM sampling (Song, Meng, and Ermon 2021) to improve sampling efficiency. For the model backbone, we use a U-Net constructed primarily from 1D convolutional layers, with parameters shared across diffusion time steps. Additional implementation details are provided in the appendix.

Order Generator The order generator consists of a simulated exchange and a meta agent. The simulated exchange replicates the double-auction market protocol on which the majority of financial markets operate. It facilitates agent-market interactions and provides the foundation for producing realistic order-level market dynamics. The meta agent represents the aggregate behavior of all traders in the generated market, functioning as a world agent. Unlike prior work that employs learned agents as world agents, our meta agent incorporates financial economic priors and is guided by the meta controller.

Specifically, the meta agent generates orders through a stochastic process, whose key parameters are supplied by the meta controller. For each trading minute t , the meta agent “wakes up” after a time interval δ_i drawn from an exponential distribution $f(\delta_i; \lambda_t) = \lambda_t e^{-\lambda_t \delta_i}$, where i indexes the total number of wake-ups within a trading day and λ_t is the arrival rate provided by the meta controller. Upon each wake-up, the meta agent instantiates an actor agent \mathcal{A}_i from a family of heterogeneous agents (Chiarella, Iori, and Perello 2009). The agent makes decisions by optimizing a CARA utility function based on current market observations. The order-generation procedure proceeds as follows:

- **Initialization.** The actor agent is initialized with a random holding position S , the corresponding cash balance C , and random component weights g_f, g_c, g_n for its three heterogeneous components, namely fundamental, chartist, and noise. These weights are drawn independently from exponential distributions whose expected values follow a ratio of 10:1.5:1.
- **Return estimation.** The actor agent estimates the objective future return \hat{r} as a weighted average of the fundamental, chartist, and noise components. The fundamental component is given by the return r_t determined by the meta controller, the chartist component is the historical average return \bar{r} obtained from the simulated exchange, and the noise component is a small Gaussian perturbation r_σ . Accordingly, $\hat{r} = \frac{g_f r_t + g_c \bar{r} + g_n r_\sigma}{g_f + g_c + g_n}$.
- **Holding optimization.** Given the estimated return, the actor agent computes a future price estimate $\hat{p}_t = p_t \exp(\hat{r})$. By optimizing the CARA utility over future wealth (Chiarella, Iori, and Perello 2009), the actor derives its demand function $u(p) = \frac{\ln(\hat{p}_t/p)}{aVp}$, where a is the risk-aversion coefficient and V denotes historical price volatility. The actor then identifies the lowest acceptable price p_l satisfying $p_l(u(p_l) - S) = C$, ensuring affordability under the current cash balance.
- **Order sampling.** The actor agent samples an order price uniformly between the lowest acceptable price and the estimated price, $p_i \sim \mathcal{U}(p_l, \hat{p}_t)$. It then determines the order volume as $q_i = u(p_i) - S$ and the order type as $o_i = \text{sign}(q_i)$, where $o_i = 1$ indicates a buy order and $o_i = 0$ indicates a sell order.

The resulting order is recorded as $\mathbf{o}_i = (t_i, p_i, q_i, o_i) \sim p(\mathbf{o} \mid r_t, \lambda_t)$, where $t_i = \sum_{j=1}^i \delta_j$. Pseudocode for the order-generation procedure is provided in Algorithm 2 in the appendix.




Figure 3: Aggregated price curves (first row) and distributions of targeted indicators (second row). In the first row, each curve represents the price trajectory derived from one day of generated order flow. In the second row, each colored density corresponds to the distribution of a targeted indicator computed from the generation results.

		A-Main					ChiNext				
Target	Method	Lower	Low	Medium	High	Higher	Lower	Low	Medium	High	Higher
Return	No Control	1.443	0.583	0.529	0.813	2.337	0.979	0.684	0.992	1.718	3.923
	Discrete	1.055	0.494	0.228	0.429	0.664	1.285	0.807	0.243	0.413	0.869
	Continuous	0.206	0.178	0.161	0.184	0.212	0.584	0.539	0.342	0.449	0.840
Amplitude	No Control	0.521	0.268	0.268	0.699	3.298	1.130	0.638	0.427	0.608	2.763
	Discrete	0.049	0.088	0.309	0.502	0.930	0.057	0.134	0.346	0.523	0.963
	Continuous	0.054	0.076	0.149	0.247	0.348	0.110	0.116	0.255	0.437	0.973
Volatility	No Control	0.021	0.115	0.431	1.209	4.288	0.029	0.246	0.713	1.737	5.221
	Discrete	0.016	0.123	0.383	0.890	2.393	0.029	0.188	0.481	0.948	2.257
	Continuous	0.011	0.104	0.318	0.774	2.389	0.028	0.178	0.473	1.016	2.631

Table 1: Mean squared error (MSE) between the targeted indicators and the aggregated statistics computed from the generated order flow. Best results are highlighted in bold.

Order generation terminates at t_{\max} when the next wake-up time t_i would exceed the trading hours of the day. The final generated order flow is recorded as

$$\tilde{\mathcal{O}} = \{\mathcal{o}_1, \dots, \mathcal{o}_{\max}\} \sim p(\mathcal{O} | \tilde{\mathbf{x}}), \quad (9)$$

where $\tilde{\mathbf{x}} = \{r, \lambda\}$ is generated by the meta controller.

Experiments

In this section, we describe the experimental setup and present results on real-world datasets to evaluate both the controllability and fidelity of DigMA. We further demonstrate the usefulness of DigMA as a generative environment for a high-frequency trading reinforcement learning task and analyze its computational efficiency.

Dataset and Model Configurations

We conduct experiments on two tick-by-tick order-flow datasets from global markets: *A-Main* and *ChiNext* from the Chinese stock market. For each dataset, we use 5,000 samples for validation and 5,000 samples for testing, with the remaining samples used for training. Additional details on

dataset preprocessing, along with links to the code and illustrative examples, are provided in the appendix.

We train the diffusion model on each dataset for 10 epochs with 200 diffusion steps using the AdamW optimizer. The mini-batch size is set to 256, and the learning rate is 1×10^{-5} . For both discrete and continuous control models, conditioning information is randomly dropped with probability 0.5 during training. Other implementation details are provided in the appendix.

Evaluation on Controlling Financial Market Generation

In this experiment, we evaluate DigMA’s ability to perform controllable market generation. To enable conditioning on target scenarios, we train DigMA using three indicators: *return*, *amplitude*, and *volatility*, each capturing a broad range of market conditions. For each indicator, we first obtain its empirical distribution from the real-world order-flow dataset. We then partition the values into five percentile-based bins, representing *lower*, *low*, *mid*, *high*, and *higher* value cases for the scenario characterized by the




Figure 4: Comparison of stylized facts distribution across baselines. The x-axis is the stylized facts and the y-axis is the density.

	A-Main				ChiNext			
Model	MinR	RetAC	VolC	OIR	MinR	RetAC	VolC	OIR
RFD	1.198	5.010	0.839	0.015	0.272	2.987	0.691	0.022
RMSC	2.640	10.170	1.237	0.563	1.371	7.461	0.668	0.588
LOGGAN	0.151	1.903	1.101	0.309	0.135	1.711	0.507	0.282
DigMA	0.084	2.781	0.273	0.009	0.079	1.997	0.218	0.009

Table 2: K-L divergence of stylized facts distribution between real and simulated order flow.

corresponding indicator. DigMA with a discrete control encoder is trained using these bin labels as categorical conditions. For the continuous control encoder, we use normalized numerical indicator values as inputs. During testing, discrete models take the corresponding class labels as control targets, while continuous models use the median indicator value of each bin computed from real samples.

Table 1 reports the **mean squared error (MSE)** between each control target (i.e., bin median) and the corresponding indicator computed from the generated order flow, for DigMA with both discrete and continuous control encoders. Results are averaged over three runs with different random seeds. The “No Control” baseline is a DigMA variant in which the conditioning mechanism of the meta-controller is removed, resulting in generation that is independent of the target scenario. As shown in Table 1, DigMA achieves consistently low errors across all indicators, whereas the uncontrolled baseline produces either random or scenario-insensitive outputs.

Figure 3 visualizes the mid-price series for controlled generation samples across different scenarios, together with the distributions of indicators computed from 200 independent runs per scenario. The price trajectories reflect the intended market conditions, and the indicator distributions shift appropriately in response to the control targets. These results demonstrate that DigMA effectively enables controllable financial market generation.

Evaluation on Generation Fidelity

We evaluate the generation fidelity of DigMA and compare it with both rule-based and learning-based baselines:

- **RFD** (Vyetenko et al. 2020) is a market simulation configuration featuring heterogeneous agents, including 1 market maker, 25 momentum traders, 100 value traders, and 5,000 noise traders with random fundamentals.
- **RMSC** (Amrouni et al. 2021) is the reference market simulation configuration introduced in ABIDES-gym. It

includes all agents in RFD, along with an additional percentage-of-volume (POV) agent that provides extra liquidity to the simulated market.

- **LOGGAN** (Coletta et al. 2023) is a conditional Wasserstein GAN with gradient penalty, trained to generate the next order conditioned on market history.

We assess fidelity by examining the distributional discrepancies of several canonical “stylized facts,” which are key statistical properties of asset returns and order-book dynamics, between real and simulated markets. These statistics capture widely studied characteristics of financial market microstructure:

- **Minutely Log Returns (MinR)** are the logarithmic differences between consecutive minute-level prices.
- **Return Auto-correlation (RetAC)** is the linear auto-correlation between the return series and its lagged values. Empirical studies on real market data show that returns exhibit little to no auto-correlation at short lags.
- **Volatility Clustering (VolC)** is the linear auto-correlation of squared returns and their lagged values. It reflects the empirical observation that periods of high volatility tend to cluster over time.
- **Order Imbalance Ratio (OIR)** is the proportional volume difference between the best bid and best ask, capturing directional trading tendencies of market participants.

Additional details are provided in the appendix.

We report results using the unconditional variant of DigMA to isolate the effect of controllability from fidelity. Figure 4 illustrates the distributions of these statistics, where the real-market distributions are shown as *Real* using a solid golden line. Across all metrics, DigMA more closely matches the real distributions than competing methods. To quantify these discrepancies, we compute the **Kullback-Leibler (K-L) divergence** between the real and simulated statistics. Table 2 summarizes the results, demonstrating that

Environment	Ret(%)(\uparrow)	Vol(\downarrow)	SR(\uparrow)	MDD(%)(\uparrow)
Replay	0.009 \pm 0.043	0.413 \pm 0.090	0.014 \pm 0.008	-1.133 \pm 0.173
RFD	0.000 \pm 0.008	0.159 \pm 0.094	0.011 \pm 0.029	-0.803 \pm 0.327
DigMA-c	0.015 \pm 0.023	0.147 \pm 0.151	0.006 \pm 0.066	-0.715 \pm 0.464
DigMA	0.029 \pm 0.019	0.411 \pm 0.121	0.049 \pm 0.031	-1.313 \pm 0.156

Table 3: Average out-of-sample test results (in percentage). Best results are highlighted with bold face.

DigMA achieves the lowest K-L divergence on most metrics. While LOBGAN attains the lowest divergence on Re-tAC due to its auto-regressive order generation, DigMA substantially outperforms LOBGAN on the remaining stylized facts by a large margin. Overall, these results indicate that DigMA achieves superior fidelity in market simulation, effectively capturing realistic market dynamics.

Evaluation on High-Frequency Trading with Reinforcement Learning

We evaluate the usefulness of DigMA as a training environment for reinforcement learning (RL) agents in a high-frequency trading task.

Settings We train a trading agent in simulated market environments using the A2C algorithm, with the objective of optimizing high-frequency trading performance. The agent takes a discrete action every 10 seconds. The action space includes buying or selling at any of the best five price levels with an integer volume between 1 and 10 units, as well as an option to take no action. The observation space consists of price changes over the past twenty seconds, ten-level bid-ask price-volume pairs, and account information including capital, position, and cash.

Each agent is trained in environments generated by one of the following methods: historical replay, RFD, DigMA, and a variant of DigMA without meta-controller conditioning (DigMA-c). Each training run lasts 200 episodes, with each episode corresponding to a full trading day containing 1,440 decision steps over four trading hours. After training, agents are evaluated for 50 episodes in an environment that replays out-of-sample real market data. This evaluation is repeated three times using three non-overlapping out-of-sample periods to assess robustness. All RL models are trained with identical hyperparameters to ensure fairness.

We evaluate the performance of the RL trading task using four metrics. **Daily return (Ret)** is defined as the mean return across episodes and is used to assess profitability. **Daily volatility (Vol)** is defined as the standard deviation of daily returns and is used to assess risk. **Sharpe ratio (SR)** is defined as the ratio of daily return to volatility, reflecting the return-risk trade-off. **Maximum drawdown (MDD)** is defined as the largest intraday decline in cumulative profit during testing and is used to assess robustness under extreme market conditions. For daily return and Sharpe ratio, higher values indicate better performance, whereas for daily volatility and maximum drawdown, lower values are preferred.

Results Table 3 reports the average numerical results of the high-frequency trading task. The trading agent trained in the DigMA-generated environment achieves the high-

Model	Time(ms)/Order
RFD	0.049
RMSC	0.075
LOBGAN	1.710
DigMA	0.017

Table 4: Comparison on computational efficiency.

est daily return and Sharpe ratio among all baselines. The agent trained with DigMA-c attains the lowest daily volatility and maximum drawdown, indicating the emergence of a more conservative trading strategy. These results demonstrate that DigMA provides a more effective environment for policy learning and enables the trading agent to learn a better policy. Furthermore, the performance differences between DigMA and DigMA-c highlight that enabling controllability in the simulated training environment influences the decision-making preferences of RL trading agents.

Analysis of Computational Efficiency

In this section, we compare the computational efficiency of DigMA with all baseline models. We evaluate efficiency using the latency of financial market generation, defined as the average time required to generate a single order. As shown in Table 4, DigMA achieves the fastest order generation speed among all methods, requiring only approximately 0.017 milliseconds per order, which makes it suitable for real-time and latency-critical applications. RFD and RMSC exhibit slightly higher latency due to their designs, which require querying the virtual exchange at each decision step. LOBGAN is approximately 100 times slower than DigMA, owing to its recurrent neural network architecture and autoregressive generation process. Overall, DigMA provides the highest computational efficiency for financial market generation.

Conclusion and Future Work

In this paper, we formulate the problem of controllable financial market generation and propose the *Diffusion Guided Meta Agent* (DigMA) model to address it. Specifically, we employ a diffusion model to capture the dynamics of market states, which are represented by time-evolving distribution parameters of mid-price return rates and order arrival rates, and we define a meta agent with financial economic priors to generate orders from the corresponding distributions. Extensive experimental results demonstrate that DigMA achieves superior controllability and generation fidelity. While this work focuses on generating the order flow of a single stock at each time step, future work will consider the correlations among multiple assets to generate more realistic markets.

Acknowledgments

This work is supported by NSFC Project (No.62192783, No.12326615).

References

- Amrouni, S.; Moulin, A.; Vann, J.; Vyetenko, S.; Balch, T.; and Veloso, M. 2021. ABIDES-gym: gym environments for multi-agent discrete event simulation and application to financial markets. In *Proceedings of the ACM International Conference on AI in Finance*.
- Aversa, M.; Nobis, G.; Hägele, M.; Standvoss, K.; Chirica, M.; Murray-Smith, R.; Alaa, A. M.; Ruff, L.; Ivanova, D.; Samek, W.; Klauschen, F.; Sanguinetti, B.; and Oala, L. 2023. DiffInfinite: Large Mask-Image Synthesis via Parallel Random Patch Diffusion in Histopathology. In *Advances in Neural Information Processing Systems*.
- Brooks, T.; Peebles, B.; Holmes, C.; DePue, W.; Guo, Y.; Jing, L.; Schnurr, D.; Taylor, J.; Luhman, T.; Luhman, E.; Ng, C.; Wang, R.; and Ramesh, A. 2024. Video generation models as world simulators.
- Brown, T. B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D. M.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*.
- Byrd, D.; Hybinette, M.; and Balch, T. H. 2020. ABIDES: Towards High-Fidelity Multi-Agent Market Simulation. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*.
- Chiarella, C.; and Iori, G. 2002. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2(5): 346.
- Chiarella, C.; Iori, G.; and Perello, J. 2009. The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows. *Journal of Economic Dynamics and Control*, 33(3): 525–537.
- Chowdhery, A.; Narang, S.; Devlin, J.; Bosma, M.; Mishra, G.; Roberts, A.; Barham, P.; Chung, H. W.; Sutton, C.; Gehrmann, S.; et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240): 1–113.
- Coletta, A.; Jerome, J.; Savani, R.; and Vyetenko, S. 2023. Conditional Generators for Limit Order Book Environments: Explainability, Challenges, and Robustness. *arXiv preprint*, arXiv:2306.12806.
- Coletta, A.; Moulin, A.; Vyetenko, S.; and Balch, T. 2022. Learning to simulate realistic limit order book markets from data as a World Agent. In *Proceedings of the ACM International Conference on AI in Finance*.
- Coletta, A.; Prata, M.; Conti, M.; Mercanti, E.; Bartolini, N.; Moulin, A.; Vyetenko, S.; and Balch, T. 2021. Towards realistic market simulations: a generative adversarial networks approach. In *Proceedings of the ACM International Conference on AI in Finance*.
- Cont, R. 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2): 223.
- Copet, J.; Kreuk, F.; Gat, I.; Remez, T.; Kant, D.; Synnaeve, G.; Adi, Y.; and Défossez, A. 2024. Simple and controllable music generation. *Advances in Neural Information Processing Systems*.
- Deng, B.; Xu, C.; Li, H.; Huang, Y.-h.; Hou, M.; and Bian, J. 2025. Tardiff: Target-oriented diffusion guidance for synthetic electronic health record time series generation. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.
- Dhariwal, P.; and Nichol, A. Q. 2021. Diffusion Models Beat GANs on Image Synthesis. In *Advances in Neural Information Processing Systems*.
- Fan, X.; Wu, Y.; Xu, C.; Huang, Y.; Liu, W.; and Bian, J. 2024. MG-TSD: Multi-Granularity Time Series Diffusion Models with Guided Learning Process. In *International Conference on Learning Representations*.
- Fürutter, F.; Muñoz-Gil, G.; and Briegel, H. J. 2024. Quantum circuit synthesis with diffusion models. *Nature Machine Intelligence*, 6(5): 515–524.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems*.
- Guo, J.; Wang, S.; Ni, L. M.; and Shum, H. 2023. Quant 4.0: Engineering Quantitative Investment with Automated, Explainable and Knowledge-driven Artificial Intelligence. *arXiv preprint*, arXiv:2301.04020.
- Guo, Y.; Yang, C.; Rao, A.; Liang, Z.; Wang, Y.; Qiao, Y.; Agrawala, M.; Lin, D.; and Dai, B. 2024. AnimateDiff: Animate Your Personalized Text-to-Image Diffusion Models without Specific Tuning. *International Conference on Learning Representations*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*.
- Ho, J.; and Salimans, T. 2022. Classifier-Free Diffusion Guidance. *arXiv preprint*, arXiv:2207.12598.
- Huang, Y.-H.; Xu, C.; Wu, Y.; Li, W.-J.; and Bian, J. 2025. Timedp: Learning to generate multi-domain time series with domain prompts. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Kollovich, M.; Ansari, A. F.; Bohlke-Schneider, M.; Zschiegner, J.; Wang, H.; and Wang, Y. B. 2023. Predict, Refine, Synthesize: Self-Guiding Diffusion Models for Probabilistic Time Series Forecasting. In *Advances in Neural Information Processing Systems*.
- Kong, Z.; Ping, W.; Huang, J.; Zhao, K.; and Catanzaro, B. 2021. DiffWave: A Versatile Diffusion Model for Audio Synthesis. In *International Conference on Learning Representations*.
- Li, H.; Huang, Y.-H.; Xu, C.; Schlegel, V.; Jiang, R.; Batista-Navarro, R.; Nenadic, G.; and Bian, J. 2025a. BRIDGE:

- Bootstrapping Text to Control Time-Series Generation via Multi-Agent Iterative Optimization and Diffusion Modeling. In *Proceedings of the International Conference on Machine Learning*.
- Li, J.; Liu, Y.; Liu, W.; Fang, S.; Wang, L.; Xu, C.; and Bian, J. 2025b. MarS: a Financial Market Simulation Engine Powered by Generative Foundation Model. In *International Conference on Learning Representations*.
- Li, J.; Wang, X.; Lin, Y.; Sinha, A.; and Wellman, M. P. 2020. Generating Realistic Stock Market Order Streams. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. In *Proceedings of the International Conference on Machine Learning*.
- Lu, J.; Clark, C.; Lee, S.; Zhang, Z.; Khosla, S.; Marten, R.; Hoiem, D.; and Kembhavi, A. 2024. Unified-IO 2: Scaling Autoregressive Multimodal Models with Vision Language Audio and Action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Lux, T.; and Marchesi, M. 1999. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719): 498–500.
- M. Bran, A.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; and Schwaller, P. 2024. Augmenting large language models with chemistry tools. *Nature Machine Intelligence*, 6(5): 525–535.
- Mizuta, T. 2020. An agent-based model for designing a financial market that works well. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*.
- Palmer, R.; Brian Arthur, W.; Holland, J. H.; LeBaron, B.; and Tayler, P. 1994. Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena*, 75(1): 264–274.
- Raberto, M.; Cincotti, S.; Focardi, S. M.; and Marchesi, M. 2001. Agent-based simulation of a financial market. *Physica A: Statistical Mechanics and its Applications*, 299(1): 319–327.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Singhal, K.; Azizi, S.; Tu, T.; Mahdavi, S. S.; Wei, J.; Chung, H. W.; Scales, N.; Tanwani, A.; Cole-Lewis, H.; Pfohl, S.; et al. 2023. Large language models encode clinical knowledge. *Nature*, 620(7972): 172–180.
- Skinnider, M. A.; Wang, F.; Pasin, D.; Greiner, R.; Foster, L. J.; Dalsgaard, P. W.; and Wishart, D. S. 2021. A deep generative model enables automated structure elucidation of novel psychoactive substances. *Nature Machine Intelligence*, 3(11): 973–984.
- Sohl-Dickstein, J.; Weiss, E. A.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the International Conference on Machine Learning*.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- Song, J.; Zhang, Q.; Yin, H.; Mardani, M.; Liu, M.; Kautz, J.; Chen, Y.; and Vahdat, A. 2023. Loss-Guided Diffusion Models for Plug-and-Play Controllable Generation. In *Proceedings of the International Conference on Machine Learning*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *International Conference on Learning Representations*.
- Tu, T.; Azizi, S.; Driess, D.; Schaekermann, M.; Amin, M.; Chang, P.-C.; Carroll, A.; Lau, C.; Tanno, R.; Ktena, I.; et al. 2024. Towards generalist biomedical AI. *NEJM AI*, 1(3): AIoa2300138.
- Vaswani, A. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Vyetrenko, S.; Byrd, D.; Petosa, N.; Mahfouz, M.; Dervovic, D.; Veloso, M.; and Balch, T. 2020. Get real: realism metrics for robust limit order book market simulations. In *Proceedings of the ACM International Conference on AI in Finance*.
- Wang, J.; Hsieh, C.-Y.; Wang, M.; Wang, X.; Wu, Z.; Jiang, D.; Liao, B.; Zhang, X.; Yang, B.; He, Q.; et al. 2021. Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning. *Nature Machine Intelligence*, 3(10): 914–922.
- Wang, X.; and Wellman, M. P. 2017. Spoofing the Limit Order Book: An Agent-Based Model. In *Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems*.
- Yuan, X.; and Qiao, Y. 2024. Diffusion-TS: Interpretable Diffusion for General Time Series Generation. In *International Conference on Learning Representations*.