

CoFinDiff: Controllable Financial Diffusion Model for Time Series Generation

Yuki Tanaka¹, Ryuji Hashimoto¹, Takehiro Takayanagi¹, Zhe Piao²,
Yuri Murayama¹ and Kiyoshi Izumi¹

¹The University of Tokyo

²Nomura Holdings, Inc.

{yuki-no1-0511, hashimoto-ryuji419, takayanagi-takehiro590}@g.ecc.u-tokyo.ac.jp
zhe.piao@nomura.com, {murayama.yuri, izumi}@sys.t.u-tokyo.ac.jp

Abstract

The generation of synthetic financial data is a critical technology in the financial domain, addressing challenges posed by limited data availability. Traditionally, statistical models have been employed to generate synthetic data. However, these models fail to capture the stylized facts commonly observed in financial data, limiting their practical applicability. Recently, machine learning models have been introduced to address the limitations of statistical models; however, controlling synthetic data generation remains challenging. We propose CoFinDiff (**Controllable Financial Diffusion** model), a synthetic financial data generation model based on conditional diffusion models that accept conditions about the synthetic time series. By incorporating conditions derived from price data into the conditional diffusion model via cross-attention, CoFinDiff learns the relationships between the conditions and the data, generating synthetic data that align with arbitrary conditions. Experimental results demonstrate that: (i) synthetic data generated by CoFinDiff capture stylized facts; (ii) the generated data accurately meet specified conditions for trends and volatility; (iii) the diversity of the generated data surpasses that of the baseline models; and (iv) models trained on CoFinDiff-generated data achieve improved performance in deep hedging task.

1 Introduction

Synthetic financial data generation is a critical technology with diverse applications in financial markets. For instance, automated trading systems, a major application of machine learning (ML) in the financial domain, require diverse data sets to be well-prepared for a variety of future scenarios. However, financial data such as price series accumulate gradually over time, which limits the availability of data for these scenarios [Assefa *et al.*, 2020]. As a result, to leverage ML-based automated trading technology, synthetic financial data generation has gained considerable research attention.

Traditionally, synthetic financial data generation has relied on statistical models like Wiener processes, ARCH, and

GARCH [Engle, 1982; Bollerslev, 1986]. However, these models lack the expressive power to capture all stylized facts¹ [Malmsten and Teräsvirta, 2010], causing synthetic data to diverge from reality.

To overcome these limitations, recent research has turned to deep generative models. The high expressiveness of deep generative models enables the generation of data that satisfy stylized facts, as demonstrated by examples such as variational autoencoders (VAEs) [Kingma and Welling, 2014], generative adversarial networks (GANs) [Goodfellow *et al.*, 2014], and diffusion models [Ho *et al.*, 2020].

While successful in realistic data generation, previous research has largely overlooked the controllability of synthetic data generation in finance, despite its critical role in financial applications. Financial markets exhibit fat-tail distributions [Vyetenko *et al.*, 2021], which may trigger extreme events such as flash crashes.² These rare events carry significant risks, yet their scarcity means that relevant data are extremely limited. Therefore, addressing data scarcity in finance necessitates the conditional generation of specific events, including extreme events. Recent studies have explored conditional GANs [Mirza and Osindero, 2014; Ni *et al.*, 2021; Li *et al.*, 2020; Wiese *et al.*, 2020] to generate synthetic financial data conditioned on historical data. While effective at capturing historical trends, this approach restricts controllability because it does not allow for the direct specification of unique market conditions. Thus, no studies have been conducted on conditional generation methods that directly specify data conditions (e.g., high volatility accompanied by a very sharp downtrend).

Aiming to promote the use of ML technologies in finance by generating synthetic financial data that capture stylized facts and offer controllability, this study proposes CoFinDiff, a conditional synthetic financial data generation model using conditional diffusion models. CoFinDiff transforms stock price series into images using the Haar wavelet [Ramsey *et al.*, 1995] and employs conditional diffusion models for data learning and generation. By incorporating trend and realized volatility, calculated from the stock price series as conditions

¹Stylized facts refer to the commonly observed features of financial data across various asset classes, such as fat tails and volatility clustering.

²A flash crash is defined as a sudden, severe drop in asset prices that recovers rapidly [Kirilenko *et al.*, 2017].

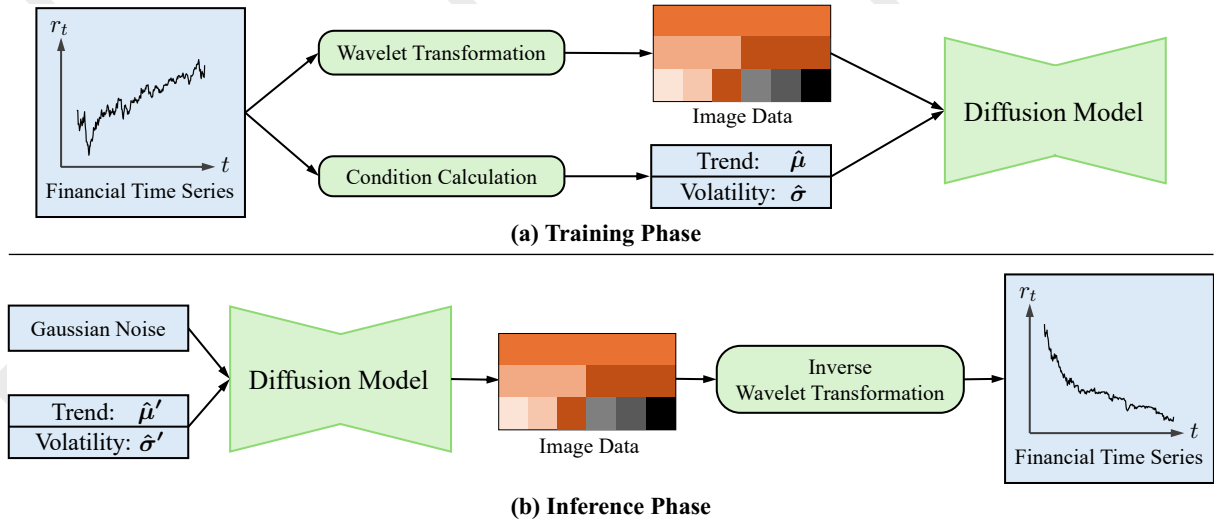


Figure 1: CoFinDiff overview. (a) In the training phase, the diffusion model learns to reproduce input logarithmic return series (Sec. 3.2). The input data is converted into images using the Haar wavelet transformation, and the trend and realized volatility are computed to serve as conditions fed into the cross-attention mechanism of the diffusion model (Sec. 3.1). (b) In the inference phase, CoFinDiff accepts arbitrary conditions using reversed wavelet transformation (Sec. 3.3).

during the learning and generation processes, CoFinDiff can generate stock price series with arbitrary trends and realized volatility. Figure 1 illustrates the overview of our proposed model.

In our experiments, we assess the following four research questions.

RQ 1: Do the synthetic data generated by CoFinDiff satisfy the stylized facts?

RQ 2: Can CoFinDiff generate synthetic data following conditions input into a model?

RQ 3: Are the synthetic data generated by CoFinDiff well diversified even under the trend and volatility constraints?

RQ 4: Are the synthetic data generated by CoFinDiff effective for downstream task?

The experimental results suggested that CoFinDiff (i) generates data that aligns with stylized facts including fat tail and volatility clustering; (ii) accurately adheres to specified conditions; (iii) generates more diverse synthetic data under fixed conditions than baselines including conditional GANs; and (iv) improves the performance of deep hedging models by providing abundant synthetic data, especially for scenarios that are rarely included in real data.

2 Related Work

Synthetic financial data have emerged as a promising tool for accelerating ML applications in finance, particularly in automated trading [Assefa *et al.*, 2020]. Various synthetic data generation methods have been developed to mitigate the limitations of real-world data availability. In particular, data-driven techniques based on deep learning have gained traction. For instance, VAEs [Kingma and Welling, 2014] and GANs [Goodfellow *et al.*, 2014] have been applied to generate synthetic financial time series, demonstrating im-

proved performance in price prediction models [Dogariu *et al.*, 2022]. GANs have been applied to generating correlated stock prices [Masi *et al.*, 2023], synthesizing price series for deep hedging [Hirano *et al.*, 2023], modeling continuous time series [Ni *et al.*, 2021], and augmenting training data for reinforcement learning-based trading models [Liu *et al.*, 2022]. Recently, diffusion models have shown promise in generating synthetic financial tabular data [Sattarov *et al.*, 2023]. Although both statistical and deep learning-based models effectively leverage real data, no generation method exists for synthetic financial data generation that captures stylized facts and enables controllability.

The primary application of synthetic financial data is deep hedging, a type of automated trading system. Deep hedging tasks, which are one of the primary applications in automated trading systems, represent a major application domain for synthetic financial data. Deep hedging tasks are major application domains of synthetic financial data. In [Buehler *et al.*, 2019], synthetic data generated via the GJR-GARCH model was employed to train a model for constructing a hedging portfolio. [Mikkilä and Kanninen, 2023] argues that training hedge strategies using artificial data based on specific statistical models is ineffective, and that training with real data yields more effective results [Mikkilä and Kanninen, 2023]. Additionally, [Hirano *et al.*, 2023] employed deep generative models to learn unnecessary hedge strategies for asset price processes, achieving promising results though hedge performance did not necessarily improve.

3 CoFinDiff

3.1 Input Data Processing

Input data processing involves setting appropriate conditions that characterize the price series and transforming data into a

suitable format for input to the diffusion model.

Calculation of Conditions Trend and realized volatility are considered as conditions characterizing the price series data, denoted as p_t , $t \in \{0, \dots, T\}$, where p_t represents the price of time t , and T denotes terminal time. By controlling the two elements, a wide range of scenarios are simulated. The trend $\hat{\mu}$ is defined as the rate of change in price. Realized volatility $\hat{\sigma}$ measures price fluctuations. Both are calculated using the logarithmic returns of price series $r_t = \log p_t - \log p_{t-1}$, $t \in \{1, \dots, T\}$ as $\hat{\mu} = \sum_{t=1}^T r_t$, and $\hat{\sigma} = \sqrt{\sum_{t=1}^T r_t^2}$. In experiments, we multiply logarithmic returns by 100 before calculating each condition to rescale appropriately.

Wavelet Transformation Diffusion models, originally designed for image processing, require an additional processing step for time series. Initially, the logarithmic returns of the price series r_t are standardized to \tilde{r}_t with a zero mean and unit variance. Subsequently, Haar wavelet transformation [Ramsey *et al.*, 1995] is applied to the standardized returns \tilde{r}_t , $t \in \{1, \dots, T\}$, and the resulting coefficients are embedded in a rectangular image format. These image representations serve as inputs to the diffusion model.

3.2 Training Phase

During training, a conditional diffusion model is used to generate price series data under specified conditions.

To ensure the model accurately learns the relationship between input data and conditions, cross-attention [Rombach *et al.*, 2022] within the model is used for conditioning. Trend and realized volatility undergo affine transformations and convolutional processing to generate the key and value for cross-attention. This processing within the attention mechanism of the conditional diffusion model enables the model to correctly learn the relationship between conditions and input data and generate corresponding data when presented with unseen conditions during inference.

3.3 Inference Phase

During the inference phase, the model trained in the training phase is used to generate data aligning with specified conditions, denoted as $\hat{\mu}', \hat{\sigma}'$. Data generation involves two major stages. In the first stage, conditions are given as input to the conditional diffusion model for generating images based on specified conditions. In the second stage, the generated images are transformed back into time series data. The process is enabled by the reversibility of the Haar wavelet transformation employed to convert time series data into images during training. Leveraging this property, the inverse wavelet transformation effectively reconstructs time series data from the generated images. Thus, the process enables the generation of price series that align with any given conditions.

4 Experimental Design

4.1 Experimental Settings

Dataset Collection This study used 1-minute FLEX full historical data provided by Japan Exchange Group.

The trading hours per day was 5 hours, with data spanning from January 1, 2015, to December 31, 2021. Missing data

Ticker	Company Name
3407	Asahi Kasei Corporation
4188	Mitsubishi Chemical Group Corporation
4568	Daiichi Sankyo Biotech Co., Ltd.
5020	Eneos Corporation
6502	Toshiba Corporation
6758	Sony Group Corporation
7203	Toyota Motor Corporation
7550	Zensho Holdings Co., Ltd.
8306	Mitsubishi UFJ Financial Group, Inc.
9202	ANA Holdings Inc.
9437	NTT Docomo, Inc.

Table 1: List of stock codes and company names.

points were added using the price from the previous time step. The stocks used for the experiment are listed in Table 1.

Training Details During training, the Adam optimizer was used. Early stopping was implemented to halt training when improvements to validation loss stopped over 100 epochs. Additionally, the number of generation steps for the diffusion model was set to 1000. The structure and parameters of the diffusion model largely followed those of the denoising diffusion probabilistic model (DDPM) [Ho *et al.*, 2020]. Given the relatively small size of the input data, the number of attention heads was set to four, while the channel size of the intermediate layers was adjusted to 20.

To capture the characteristics of extreme events, upsampling was applied to the training dataset. Data corresponding to the top 25%, 10%, and 5% of absolute trend values were replicated five times each. Upsampling ensured that extreme event data appeared multiple times within a single epoch. In financial markets, extreme events, although infrequent, entail substantial risk and consequently warrant attention.

4.2 Evaluations

Baselines

Baseline models include the GBM, GARCH model, vanilla GAN and Wasserstein GAN; in the experiment addressing RQ1, all these baselines facilitated the comparison with CoFinDiff, whereas in the experiments addressing RQ2, RQ3 and RQ4 only the GAN-based approaches (vanilla GAN and Wasserstein GAN) served as baselines.

For GBM, price path p_t , $t \in \{0, \dots, T\}$ was sampled as

$$dp_t = \mu p_t dt + \sigma p_t dW_t \quad (1)$$

dt denotes an infinitesimal time interval, dW_t denotes the corresponding increment of the Wiener process, and μ and σ denote the drift and volatility terms respectively. The GARCH(1, 1) model was used to generate return series. Return r_t and variance σ_t^2 of the model were calculated, where ϵ_t adhered to a standard normal distribution.

$$\begin{aligned} r_t &= \sigma_t \epsilon_t \\ \sigma_t^2 &= \omega + \lambda r_{t-1}^2 + \nu \sigma_{t-1}^2 \end{aligned} \quad (2)$$

where ω , λ , ν are the parameters that respectively capture the baseline variance, the effect of past squared returns, and the persistence of volatility. The structures of the vanilla GAN and Wasserstein GAN models adhere to the designs specified in [Mao *et al.*, 2017] and [Arjovsky *et al.*, 2017], respectively.

RQ 1: Evaluation of stylized facts

To answer RQ 1, an experiment was performed to verify that the synthetic data generated by CoFinDiff exhibited the property of stylized facts, traditionally observed in financial time series data [Vyetenko *et al.*, 2021]. In this study, we focus on two prominent stylized facts, fat tail and volatility clustering, which are widely recognized in financial time series [Cont, 2001].

Fat tail Fat tails, characterized by excessive probability mass in the distribution tails compared to a normal distribution, are a hallmark of financial asset returns, resulting in frequent observations of abrupt price movements. To verify the presence of fat tails in the generated data, two primary metrics were considered, Fisher’s kurtosis and the Hill index. Fisher’s kurtosis was calculated as

$$\text{Kurtosis} = \frac{T(T+1)}{(T-1)(T-2)(T-3)} \sum_{t=1}^T \left(\frac{r_t - \bar{r}}{s} \right)^4 - \frac{3(T-1)^2}{(T-2)(T-3)} \quad (3)$$

where \bar{r} and s denote the mean and standard deviation of time series r_t , respectively. Kurtosis quantified the sharpness of the distribution compared to normal distribution, with positive values indicating fat tails [Fama, 1965]. It was calculated for intraday returns, while the mean and standard deviation of the observed days were used for evaluation. The second measure involved the Hill index [Hill, 1975], calculated as

$$\xi_{(k(n),n)}^{\text{Hill}} = \left(\frac{1}{k(n)} \sum_{i=n-k(n)+1}^n \log \left(\frac{r_{(i,n)}}{r_{(n-k(n)+1,n)}} \right) \right)^{-1} \quad (4)$$

where n represents the total number of observations, $k(n)$ denotes the number of extreme observations, and $r_{i,n}$ indicates the i th order statistic of the returns sorted in increasing order.

The Hill index served as a maximum likelihood estimator for the exponent parameter of the Pareto distribution. A typical Hill index for stock markets is known to be approximately three [Gabaix *et al.*, 2006]. For this study, the Hill index was calculated using the absolute values of logarithmic returns over 1500 d, with the threshold set to top 5% of values. Drawing from empirical evidences [Lux, 1998; Gopikrishnan *et al.*, 1999; Gabaix *et al.*, 2003], we defined the stylized fact as $\xi_{(k(n),n)}^{\text{Hill}}$ in between 2.80 and 3.40.

Volatility clustering Financial markets exhibit volatility clustering, a phenomenon characterized by alternating periods of high and low price volatility [Cont, 2001]. The pattern can be identified by examining the autocorrelation of absolute returns. For this study, the autocorrelation of time series data was calculated as




Figure 2: Scatter plots of trend and realized volatility for the conditions used in training and those used in validation.

$$\rho(\tau) = \frac{\sum_{t=1}^{T-\tau} (r_t - \bar{r})(r_{t+\tau} - \bar{r})}{\sum_{t=1}^T (r_t - \bar{r})^2} \quad (5)$$

where τ is the time lag.

For comparative analysis, we used real data (Toyota Motor Corporation) and synthetic data obtained from the GBM, GARCH model, vanilla GAN, Wasserstein GAN, and CoFinDiff. Results on other stocks used in the experiments were presented in the supplementary materials. For GBM, $\mu = 0$ and $\sigma = 1$ were set. For GARCH model, typical parameters were set as $\omega = 0.1$, $\lambda = 0.1$, $\nu = 0.8$ to model stock prices [Chou, 1988]. Conditions provided to CoFinDiff and GANs were fixed, trend = 0 and realized volatility = 1. For consistency, the outputs of all methods were defined with 300 data points per day. The data generated by the GBM and CoFinDiff each covered 1500 day.

RQ 2: Controlling generation with specific conditions

During experimentation, the capacity of CoFinDiff to generate synthetic data adhering to various conditions was explored to answer RQ 2. Vanilla GAN and Wasserstein GAN were selected for comparison analysis. These models were provided with conditions not included in the training set for validation. Figure 2 shows the distribution of conditions in both real data and data for validation. The validation conditions were set on a grid corresponding to a wide range of trend and volatility values, aimed at assessing the ability to reproduce extreme market situations. Evaluation of conditional generation accuracy employs mean absolute error (MAE) as the performance metric, with method comparison achieved by computing indicators corresponding to each condition from synthetic data generated under the respective conditions and subsequently calculating the MAE.

RQ 3: Diversity of Synthetic Data

An experiment was conducted to verify whether CoFinDiff generated diverse data that satisfied given conditions to answer RQ 3. From the perspective of using synthetic data,

	Real data (7203)	GBM	GARCH (1, 1)	vanilla GAN	Wasserstein GAN	CoFinDiff
Kurtosis	6.97 (± 10.60)	0.00 (± 0.13)	0.18 (± 0.46)	5.39 (± 4.97)	4.27 (± 4.91)	5.16 (± 5.13)
Hill	3.07	5.97	5.42	3.06	3.40	2.98
Acorr (1)	0.19 (± 0.10)	0.00 (± 0.06)	0.02 (± 0.07)	0.24 (± 0.10)	0.21 (± 0.10)	0.20 (± 0.10)
Acorr (5)	0.13 (± 0.08)	0.00 (± 0.06)	0.06 (± 0.07)	0.17 (± 0.09)	0.15 (± 0.09)	0.16 (± 0.09)
Acorr (10)	0.10 (± 0.08)	0.00 (± 0.06)	0.03 (± 0.06)	0.15 (± 0.08)	0.10 (± 0.08)	0.12 (± 0.08)
Acorr (20)	0.07 (± 0.07)	0.00 (± 0.06)	0.00 (± 0.06)	0.08 (± 0.06)	0.06 (± 0.07)	0.08 (± 0.07)
Acorr (30)	0.06 (± 0.07)	0.00 (± 0.06)	0.00 (± 0.06)	0.06 (± 0.06)	0.05 (± 0.07)	0.06 (± 0.07)
fat tail	✓			✓	✓	✓
volatility clustering	✓			✓	✓	✓

Table 2: Computed statistics representing stylized facts for both real and synthetic data. Terms in parentheses correspond to the standard deviation of each statistic. Real data comprise individual stock data, while synthetic data are generated using GBM, GARCH model, GANs, and the proposed method. "Hill" refers to the Hill Index and "Acorr" denotes autocorrelation. Check marks indicate that the corresponding stylized fact is satisfied.

generating diverse data that meets the necessary properties and conditions is desirable.

For the comparative analysis, we utilized both vanilla GAN and Wasserstein GAN. We assessed the diversity of synthetic data produced by each method using dynamic time warping (DTW) and Euclidean distance. Specifically, the conditions were fixed and 200 synthetic data samples were generated for each condition. Further, the DTW and Euclidean distances for all possible pairs were calculated to determine the diversity of the synthetic data from their distributions. The fixed trend and volatility conditions were set to two combinations: (10, 50) and (-10, 50), reflecting extreme events.

RQ 4: Downstream task

Further, the effectiveness of synthetic data generated by CoFinDiff was evaluated through deep hedging tasks to answer RQ 4. The deep hedging task involves constructing a hedge portfolio for financial derivatives³ using deep learning models.

For effective hedging of financial derivatives, one must account for various factors such as transaction fees. Due to the extensive time required to manually take these factors into considerations, deep learning methods for hedging are gaining traction, leading to extensive research in deep hedging. Our study aims to enhance hedge strategies by generating training data for the deep hedging task using CoFinDiff, which can produce diverse data that satisfy stylized facts and adhere to specified conditions.

The experimental setup is as follows. European call option was used as financial derivatives. A European call option is a type of financial derivative that grants the holder the right to purchase the underlying asset at a predetermined strike price⁴ on the expiration date. The payoffs are calculated as

$$\text{European Call Option Payoff} = \max(S_T - K, 0) \quad (6)$$

where S_T denotes the asset price at maturity T , and K represents the strike price of the option. The compared asset price processes were as follows: (i) CoFinDiff, (ii) vanilla

³Derivatives are financial instruments whose value is derived from the performance of an underlying asset or group of assets.

⁴The strike price is the fixed price at which an option holder can buy or sell the underlying asset when exercising the option.

GAN, (iii) Wasserstein GAN, and (iv) real data. We selected a 5-layer perceptron for learning hedge strategies, which utilized ReLU activation function [Nair and Hinton, 2010] and layer normalization [Ba, 2016]. The conditions for generating the synthetic data were based on the trend $\hat{\mu}$ and realized volatility $\hat{\sigma}$ of the real data. Based on these two values, we defined the conditions for generating synthetic data, and for each condition, we generated data 20 times. The transaction cost rate was set at 1.00×10^{-4} . The option expiration was set to 300 minutes, consistent with previous experiments. The prediction model took as inputs the underlying asset price, time until maturity, the underlying asset's position from the previous time step, volatility, and a set of indicators (delta, gamma, theta) derived from the Black-Scholes equation [Black and Scholes, 1973]. For volatility, the real data adopted the previous day's realized volatility, whereas the generative models utilized the realized volatility given as a condition. For training using CoFinDiff and real data, data from January 1, 2015, to December 31, 2019, for all stocks listed in Table 1 were utilized, and data from 2020 were used for early stop during training and data after 2021 were employed to evaluate hedge strategies. For simplicity, all the price data used for training hedge strategies were rescaled by dividing the price at each time point by the initial price. The following two evaluation metrics were employed: (i) Entropic Risk Measure(ERM) and (ii) Conditional Value-at-Risk (CVaR).

- Entropic Risk Measure:

$$\text{ERM}_\gamma(X) = \frac{1}{\gamma} \log(\mathbb{E}[e^{-\gamma X}]) \quad (7)$$

- Conditional Value-at-Risk:

$$\text{CVaR}_\alpha(X) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_u(X) du \quad (8)$$

Here, \mathbb{E} denotes the expected value and VaR means the risk of extreme adverse outcomes, expressed as

$$\text{VaR}_u(X) = \inf\{x \mid \mathbb{P}(X \leq x) \geq u\} \quad (9)$$

where L denotes the loss random variable and u denotes the confidence level.

	Wasserstein GAN		vanilla GAN		CoFinDiff	
	(10, 50)	(-10, 50)	(10, 50)	(-10, 50)	(10, 50)	(-10, 50)
DTW	15.88(± 5.72)	14.08(± 6.38)	18.07(± 11.56)	15.91(± 8.63)	22.58(± 12.81)	16.76(± 9.32)
Euclidean	36.56(± 15.01)	25.97(± 10.89)	42.91(± 22.33)	33.35(± 15.82)	60.30(± 23.68)	52.01(± 22.97)

Table 3: Diversity of synthetic data generated under specific trend and Realized Volatility (RV) conditions for each generative method: CoFinDiff, vanilla GAN, and Wasserstein GAN. We generate 200 data under the same conditions and diversity is calculated by the average of the distances in each indicator for all combinations. $(\hat{\mu}, \hat{\sigma})$ indicates that the given condition is trend $\hat{\mu}$ and realized volatility $\hat{\sigma}$.

	vanilla GAN	Wasserstein GAN	CoFinDiff
Trend	1.83	1.34	1.02
RV	9.99	16.21	5.94

Table 4: The average absolute error of synthetic data generated under Trend and Realized Volatility (RV) conditions for each generative method: CoFinDiff, vanilla GAN, and Wasserstein GAN.

Evaluating the deep hedging task using the entire test dataset produces an average performance metric that obscures the assessment of hedging risk in individual cases. In the deep hedging task, one should examine separately those cases in which hedging plays a crucial role, namely the situations in which the absence of hedging leads to significant losses to verify the effectiveness of the hedging portfolio construction. In this context, we considered two cases: (i) The first case assumes an upward trend in the underlying asset price. In our experiment, the European call option yields a payoff only when the asset price exceeds the strike price, so a rising asset price increases risk. Thus, training and test data exhibiting an upward trend were used. (ii) The second case assumes significant fluctuations in the underlying asset price. High volatility leads to large price movements and increases the likelihood of substantial payoffs, thereby raising risk as in case (i). For this case, training and test data were selected from periods with realized volatility in the top 10% or higher. To ensure neutrality in trend prediction, each synthetic dataset was normalized to a zero trend before applying volatility filtering.

5 Experimental Results and Discussion

RQ 1: Evaluation of stylized facts

Table 2 presents the calculated results for Fisher’s kurtosis, Hill index, and autocorrelation in the absolute values of the return series. The results indicate that synthetic data generated by the three models, CoFinDiff, vanilla GAN and Wasserstein GAN, satisfy the properties of two stylized facts, fat tails and volatility clustering.

Fat tail For kurtosis, the series generated by GBM and GARCH model exhibit values near zero, while the real and synthetic data generated by GANs and proposed model show positive values. Similarly, the Hill index for GBM and GARCH model exceeds the typical value of three observed in stock prices, contrary to the value of approximately three observed for real and synthetic data generated by GANs and CoFinDiff. These results indicate that the proposed conditional diffusion model successfully captured the fat tail characteristics observed in real data.

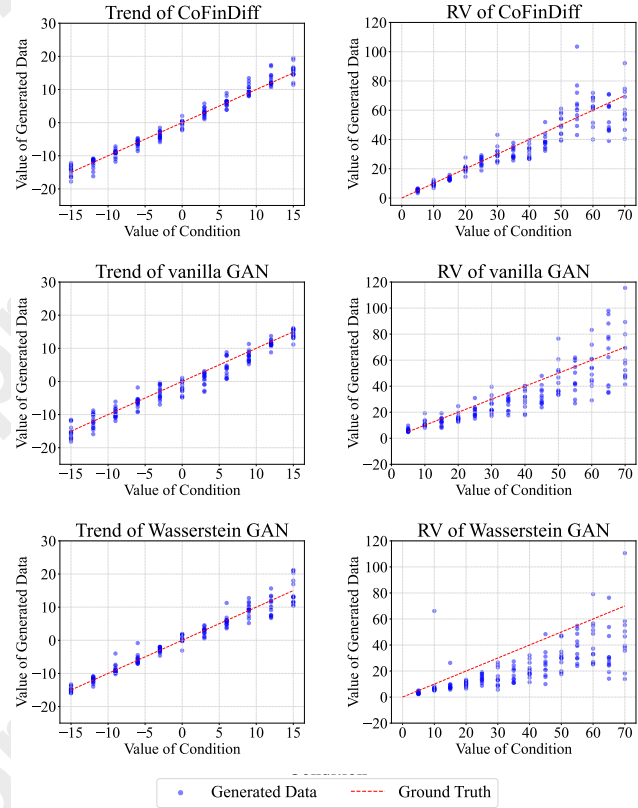


Figure 3: Scatter plots of trend and Realized Volatility (RV) conditions along with the corresponding synthetic data metrics for each generative method: CoFinDiff, vanilla GAN, and Wasserstein GAN.

Volatility clustering While GBM exhibits negligible autocorrelation, real data and data generated by GARCH and CoFinDiff demonstrate positive autocorrelation. In particular, both real data and data generated by CoFinDiff exhibit significant positive autocorrelation, confirming the presence of volatility clustering in the latter two. Thus, the proposed conditional diffusion model successfully captured the volatility clustering phenomenon in financial markets.

RQ 2: Controlling generation with specific conditions

Figure 3 shows a scatter plot illustrating the relationship between trend and realized volatility conditions provided to each model, along with corresponding values of generated data, and Table 4 shows the mean absolute error between the input conditions and corresponding values of generated data. The results indicate that CoFinDiff achieves higher

Data	Utility	Real data	vanilla GAN	Wasserstein GAN	CoFinDiff
All data	ERM($\gamma=100$)	0.00571	0.00624	<i>0.00605</i>	0.00609
All data	CVaR($\alpha=0.05$)	0.01298	0.01533	0.01321	<i>0.01319</i>
Only up trend	ERM($\gamma=100$)	<i>-0.03149</i>	-0.02189	-0.02966	-0.03772
Only up trend	CVaR($\alpha=0.05$)	<i>-0.00910</i>	0.00002	-0.00746	-0.01097
Only high volatility	ERM($\gamma=100$)	<i>0.00768</i>	0.00875	0.00790	0.00762
Only high volatility	CVaR($\alpha=0.05$)	0.01780	0.01830	0.01645	<i>0.01722</i>

Table 5: The results of training prediction models using four approaches — real data, CoFinDiff, vanilla GAN and Wasserstein GAN — and constructing and evaluating a hedging portfolio are presented. Here, ERM denotes the Entropic Risk Measure, and CVaR stands for Conditional Value-at-Risk. The best and second-best values are highlighted in bold and italic, respectively.

accuracy in conditional synthetic financial data generation than vanilla GAN and Wasserstein GAN. Although vanilla GAN exhibits notable performance under both trend conditions and realizes volatility conditions, as discussed further, the model is limited by mode collapse. This phenomenon results in the generation of highly similar data, thereby preventing the generation of useful synthetic data. Regarding trend conditions, both CoFinDiff and the Wasserstein GAN satisfy the conditions with similar high accuracy. Contrarily, for realized volatility conditions, CoFinDiff generates synthetic data with higher accuracy than the Wasserstein GAN. Specifically, when the realized volatility is high, a significant difference in accuracy between CoFinDiff and GAN is observed. High realized volatility indicates dramatically fluctuating prices, corresponding to extreme events. Therefore, CoFinDiff demonstrates particularly outstanding performance in generating synthetic data for extreme events, where data scarcity poses a serious challenge.

RQ 3: Diversity of Synthetic Data

Table 3 presents the evaluation results of diversity in synthetic data generated by each method. The results indicate that CoFinDiff achieves the highest diversity in synthetic data. Vanilla GAN demonstrates low diversity across both evaluation metrics, which is likely to be attributed to mode collapse commonly observed in GANs. Wasserstein GAN, employing the Wasserstein distance during training to prevent mode collapse, shows some mitigation of the issue. However, a significant disparity is observed between CoFinDiff and Wasserstein GAN under conditions corresponding to extreme events. Therefore, CoFinDiff is capable of generating diverse data, particularly under conditions associated with extreme events.

RQ 4: Downstream task

The results are presented in Table 5. The results indicate that evaluations employing the complete test dataset yield superior performance for models trained on actual data, while evaluations under conditions of heightened hedging necessity reveal that models trained on data generated by CoFinDiff tend to achieve the highest performance.

When all available data served as test data, the model trained on real data achieved the best performance. One explanation attributes this result to the situation in which abundant training data remains available without assuming any specific scenario; under these circumstances, real data, which best reflects the properties of the test set, proves most suit-

able. In contrast, under case (i) that assumes rising underlying asset prices and case (ii) that assumes high volatility of the underlying asset, the model trained on CoFinDiff synthetic data yielded the highest performance. One explanation attributes this outcome to the reduced amount of real training data available when assuming specific scenarios; in such cases, CoFinDiff synthetic data, which provides a sufficient sample size along with high accuracy and diversity in conditions, becomes most effective.

6 Conclusion and Future work

This study proposes a model for generating financial time series data using a conditional diffusion model, conditioned on trends and realized volatility. The primary research question involves whether it is possible to generate diverse financial time series data that satisfy stylized facts and adhere to arbitrary conditions. To investigate this, we evaluate the following aspects: (i) whether the price series generated by the diffusion model meets stylized facts such as fat tails, (ii) whether the generated price series satisfies the specified trends and realized volatility, (iii) whether generated price series exhibit diversity, and (iv) whether in the deep hedging task, a model trained on synthetic data is capable of effectively constructing a hedging portfolio. The experimental results confirm that the proposed conditional diffusion model meets all the criteria.

Although this study significantly contributes to the generation of financial time series data, substantial potential for further development remains. Future work can be directed toward two main areas.

First, the generation targets can be expanded. Although the current study generates data without being limited to specific stocks, it is possible to refine the model to generate data for particular stocks by conditioning on stock identifiers. Additionally, the model can be adapted to capture correlations between multiple stocks, enabling the generation of multi-stock price series. Expanding to generate other types of financial time series data, such as trading volumes or spreads, serves as a feasible direction.

The two avenues involve extending the conditions used in the generation process. This study utilizes only the daily price trend and realized volatility as a condition. Future work can explore a broader range of conditions to achieve more flexible conditional generation. Alternatively, incorporating natural language descriptions of conditions can further enhance the model’s flexibility in generating financial time series data.

Acknowledgments

This work was supported in part by Nomura Holdings, Inc.

References

- [Arjovsky *et al.*, 2017] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pages 214–223. PMLR, 2017.
- [Assefa *et al.*, 2020] Samuel A Assefa, Danial Dervovic, Mahmoud Mahfouz, Robert E Tillman, Prashant Reddy, and Manuela Veloso. Generating synthetic data in finance: Opportunities, challenges and pitfalls. In *Proceedings of ACM International Conference on AI in Finance*, pages 1–8, 2020.
- [Ba, 2016] Jimmy Lei Ba. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [Black and Scholes, 1973] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [Bollerslev, 1986] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.
- [Buehler *et al.*, 2019] Hans Buehler, Lukas Gonon, Josef Teichmann, and Ben Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [Chou, 1988] Ray Yeutien Chou. Volatility persistence and stock valuations: Some empirical evidence using garch. *Journal of Applied Econometrics*, 3(4):279–294, 1988.
- [Cont, 2001] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223, 2001.
- [Dogariu *et al.*, 2022] Mihai Dogariu, Liviu-Daniel Ștefan, Bogdan Andrei Boteanu, Claudiu Lamba, Bomi Kim, and Bogdan Ionescu. Generation of realistic synthetic financial time-series. *ACM Trans. Multimedia Comput. Commun. Appl.*, 18(4), 2022.
- [Engle, 1982] Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007, 1982.
- [Fama, 1965] Eugene F Fama. The behavior of stock-market prices. *The journal of Business*, 38(1):34–105, 1965.
- [Gabaix *et al.*, 2003] Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–70, 06 2003.
- [Gabaix *et al.*, 2006] Xavier Gabaix, Parameswaran Gopikrishnan, Vasiliki Plerou, and H Eugene Stanley. Institutional investors and stock market volatility. *The Quarterly Journal of Economics*, 121(2):461–504, 2006.
- [Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- [Gopikrishnan *et al.*, 1999] Parameswaran Gopikrishnan, Vasiliki Plerou, Luís Amaral, Martin Meyer, and H. Stanley. Scaling of the distribution of fluctuations of financial market indices. *Physical Review. E, Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics*, 60:5305–16, 12 1999.
- [Hill, 1975] Bruce M. Hill. A simple general approach to inference about the tail of a distribution. *The Annals of Statistics*, 3(5):1163–1174, 1975.
- [Hirano *et al.*, 2023] Masanori Hirano, Kentaro Minami, and Kentaro Imajo. Adversarial deep hedging: Learning to hedge without price process modeling. In *Proceedings of ACM International Conference on AI in Finance*, pages 19–26, 2023.
- [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.
- [Kingma and Welling, 2014] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [Kirilenko *et al.*, 2017] Andrei Kirilenko, Albert S Kyle, Mehrdad Samadi, and Tugkan Tuzun. The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3):967–998, 2017.
- [Li *et al.*, 2020] Junyi Li, Xintong Wang, Yaoyang Lin, Arunesh Sinha, and Michael Wellman. Generating realistic stock market order streams. In *Proceedings of AAAI Conference on Artificial Intelligence*, volume 34, pages 727–734, 2020.
- [Liu *et al.*, 2022] Chunli Liu, Carmine Ventre, and Maria Polukarov. Synthetic data augmentation for deep reinforcement learning in financial trading. In *Proceedings of the ACM International Conference on AI in Finance*, page 343–351, 2022.
- [Lux, 1998] Thomas Lux. The limiting extremal behaviour of speculative returns: An analysis of intra-daily data from the Frankfurt stock exchange. *IFAC Proceedings Volumes*, 31(16):29–33, 1998.
- [Malmsten and Teräsvirta, 2010] Hans Malmsten and Timo Teräsvirta. Stylized facts of financial time series and three popular models of volatility. *European Journal of Pure and Applied Mathematics*, 3(3):443–477, 2010.
- [Mao *et al.*, 2017] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of IEEE International Conference on Computer Vision*, pages 2794–2802, 2017.
- [Masi *et al.*, 2023] Giuseppe Masi, Matteo Prata, Michele Conti, Novella Bartolini, and Svitlana Vyetenko. On correlated stock market time series generation. In *Proceed-*

ings of ACM International Conference on AI in Finance, page 524–532, 2023.

- [Mikkilä and Kanninen, 2023] Oskari Mikkilä and Juho Kanninen. Empirical deep hedging. *Quantitative Finance*, 23(1):111–122, 2023.
- [Mirza and Osindero, 2014] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [Nair and Hinton, 2010] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the International Conference on Machine Learning*, pages 807–814, 2010.
- [Ni *et al.*, 2021] Hao Ni, Lukasz Szpruch, Marc Sabate-Vidales, Baoren Xiao, Magnus Wiese, and Shujian Liao. Sig-wasserstein gans for time series generation. In *Proceedings of ACM International Conference on AI in Finance*, pages 1–8, 2021.
- [Ramsey *et al.*, 1995] James B. Ramsey, Daniel Usikov, and George M. Zaslavsky. An analysis of U.S. stock price behavior using wavelets. *Fractals*, 03(02):377–389, 1995.
- [Rombach *et al.*, 2022] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [Sattarov *et al.*, 2023] Timur Sattarov, Marco Schreyer, and Damian Borth. FinDiff: Diffusion models for financial tabular data generation. In *Proceedings of ACM International Conference on AI in Finance*, page 64–72, 2023.
- [Vyetenko *et al.*, 2021] Svitlana Vyetenko, David Byrd, Nick Petosa, Mahmoud Mahfouz, Danial Dervovic, Manuela Veloso, and Tucker Balch. Get real: Realism metrics for robust limit order book market simulations. In *Proceedings of the ACM International Conference on AI in Finance*, 2021.
- [Wiese *et al.*, 2020] Magnus Wiese, Robert Knobloch, Ralf Korn, and Peter Kretschmer. Quant GANs: Deep generation of financial time series. *Quantitative Finance*, 20(9):1419–1440, 2020.