

Advanced Case Study: Market-GAN and the Evolution of Financial Market Generators

ISSN: 2832-4463



***Corresponding author:** Saqib Qamar B, Department of Intelligent Systems, KTH Royal Institute of Technology, Sweden

Submission:  July 07, 2025

Published:  September 03, 2025

Volume 4- Issue 5

How to cite this article: Saqib Qamara B* and Mohd Faizanc. Advanced Case Study: Market-GAN and the Evolution of Financial Market Generators. COJ Rob Artificial Intel. 4(5). COJRA. 000598. 2025.
DOI: [10.31031/COJRA.2025.04.000598](https://doi.org/10.31031/COJRA.2025.04.000598)

Copyright@ Saqib Qamara B, This article is distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use and redistribution provided that the original author and source are credited.

Saqib Qamara B^{1,2*} and Mohd Faizanc³

¹Department of Computing and IT, Sohar University, Oman

²KTH Royal Institute of Technology, Sweden

³Department of Electronics, Zakir Husain College of Engineering and Technology, Aligarh Muslim University, India

Abstract

Financial market simulation has evolved significantly with the advent of generative models, enabling high-fidelity synthetic data generation for forecasting, risk management, and algorithmic trading. This case study explores Market-GAN, a novel framework for controllable financial data generation with semantic context, and situates it within the broader paradigm shift toward generative AI in finance. We analyze Market-GAN's architecture, training methodology, and performance against benchmarks, while also comparing it with other state-of-the-art market generators like Sig-Wasserstein GANs and VoLGAN. The study highlights key innovations such as contextual conditioning, two-stage adversarial training, and signature-based evaluation, while addressing challenges like non-stationarity, data scarcity, and evaluation metrics. Finally, we discuss future directions, including the integration of foundation models and reinforcement learning for next-generation financial simulators.

Keywords: Market-GAN; Generative adversarial networks; Financial simulation; Synthetic data; Signature methods; Contextual generation

Introduction




Figure 1: Overview of the market dynamics.

Financial markets represent complex adaptive systems characterized by non-stationary dynamics, structural breaks, and intricate inter-dependencies across multiple time scales as shown in Figure 1. The challenge of accurately simulating these markets has become increasingly critical for three key applications: (i) risk management under stressed market conditions, (ii) development and back testing of algorithmic trading strategies, and (iii) regulatory stress testing of financial institutions. Traditional simulation approaches including Stochastic Differential Equation (SDE) models and agent-based simulations face fundamental limitations in capturing the nuanced behavior of modern financial markets [1,2]. While SDE models like Heston or SABR provide elegant mathematical frameworks, their parametric

assumptions often fail to accommodate regime shifts and extreme events. Similarly, agent-based models, despite their micro-foundational appeal, struggle with computational scalability and empirical validation (Figure 1).

The advent of generative artificial intelligence, particularly Generative Adversarial Networks (GANs) and their variants, has ushered in a new paradigm for financial market simulation [3]. These data-driven approaches offer three transformative advantages: (i) they learn directly from historical market data without restrictive parametric assumptions, (ii) they can capture non-linear dependencies and tail risk characteristics that elude traditional models, and (iii) they enable conditional generation of market scenarios based on specific macroeconomic or microstructural contexts.

This study presents a comprehensive analysis of Market-GAN [4], a groundbreaking framework that advances financial simulation through three key innovations. First, it introduces semantic context including market regimes, asset-specific characteristics, and historical patterns as control variables for conditional generation. Second, it combines the strengths of adversarial training with autoencoder-based feature extraction and supervisor networks for context preservation. Third, it implements a novel two-stage training protocol that addresses the persistent challenge of mode collapse in financial GAN applications.

We situate Market-GAN within the broader landscape of modern market generators through comparative analysis with two other significant approaches: (i) Sig-Wasserstein GANs [5], which employ path-space metrics from rough path theory to ensure temporal consistency, and (ii) VoLGAN [6], which specializes in generating arbitrage-free implied volatility surfaces for derivatives markets. Our investigation addresses five critical dimensions:

1. Architectural innovations that enable controllable generation while preserving market microstructure properties
2. Training methodologies that balance stability with model capacity
3. Comprehensive evaluation frameworks that assess both statistical fidelity and economic plausibility
4. Comparative performance against both classical and contemporary benchmarks
5. Future research directions at the intersection of generative finance and reinforcement learning

The practical implications of this research extend across multiple financial domains. For quantitative hedge funds, advanced market generators enable more robust strategy development through improved scenario analysis. For central banks and regulators, they provide tools for systemic risk assessment under hypothetical stress scenarios. For fintech innovators, they offer pathways to develop next-generation trading algorithms while addressing data privacy concerns through synthetic data generation. This manuscript is organized as follows: Section 2 reviews the evolution

of financial simulation methods and establishes the theoretical foundations. Section 3 details Market-GAN's architecture and training methodology. Section 4 presents our experimental framework and results. Section 5 provides comparative analysis with alternative approaches. Section 6 discusses limitations and future directions, concluding with implications for both academic research and financial practice.

Background and Related Work

Classical financial simulation approaches

The theoretical foundations of financial market simulation trace back to the pioneering work of Black et al. [7], whose option pricing models established geometric Brownian motion as the standard framework for asset price modeling. However, as noted by [8], these early models failed to capture several empirically observed market phenomena, prompting the development of stochastic volatility models. Heston [9] introduced a square-root diffusion process for volatility, while later extensions like the SABR model [10] improved smile fitting capabilities. For modeling discontinuous price movements, Merton's [11] jump-diffusion framework provided important theoretical advances. The ARCH/GARCH family of models, originating with Engle & Bollerslev [12,13], offered an alternative approach to volatility modeling through autoregressive conditional heteroskedasticity. As demonstrated by Andersen et al. [14], these models successfully captured volatility clustering but remained limited in their ability to model high-frequency dynamics. The limitations of purely stochastic approaches led to the development of agent-based modeling frameworks, with seminal contributions including the market microstructure models of Glosten et al. [15] and the heterogeneous agent approaches of Brock et al. [16].

Time series analysis in finance evolved through several generations, from the ARIMA methodology of Box et al. [17] to state-space approaches popularized by Harvey [18]. The application of neural networks began with relatively simple architectures as described by Refenes et al. [19], with later advances incorporating recurrent connections as in the LSTM architecture of Hochreiter et al. [20]. However, as noted by Zhang et al. [21], these early neural approaches required extensive feature engineering and struggled with financial data's non-stationarity.

Modern generative approaches in finance

The application of deep generative models to financial data represents a significant paradigm shift, building on foundational work in machine learning. Goodfellow et al. [22] introduced Generative Adversarial Networks, while Kingma et al. [23] developed the Variational Autoencoder framework. In financial time series generation, Esteban et al. [24] first adapted GANs through their RCGAN architecture, with subsequent improvements in temporal modeling coming from Yoon et al. [25] in Time GAN. Signature-based methods draw from the mathematical foundations of rough path theory developed by Lyons [26]. The application to financial modeling was pioneered by Lyons et al. [27], with recent advances including the Sig-WGAN of Ni Hao et al. [5] and neural

SDE approaches of Kidger et al. [28]. These methods provide theoretically grounded approaches to capturing path dependencies, addressing limitations noted in traditional time series models by Cont [8].

Market-GAN builds upon these foundations while introducing several key innovations. The contextual conditioning framework extends concepts from conditional GANs [29] and builds on clustering approaches similar to those used by Khandani et al. [30] for regime identification. The hybrid architecture combines elements from autoencoder-based feature learning [31] with adversarial training, while the financial constraints incorporate insights from arbitrage-free modeling. This synthesis of techniques addresses many of the limitations identified in previous financial simulation approaches while maintaining the theoretical rigor emphasized by Cont et al. [32].

Market-GAN: Architecture and Methodology

Contextual market dataset

The Market-GAN framework processes a comprehensive dataset structure that captures both price dynamics and contextual market information. Each data sample consists of price features represented as Open-High-Low-Close (OHLC) vectors, accompanied by three critical contextual components. Market dynamics are classified into distinct regimes (bull, bear, or flat markets) through an innovative clustering algorithm that analyzes price trajectories and volatility patterns as seen in Figure 1. Asset-specific context is incorporated via ticker embeddings that capture unique characteristics of individual securities. Temporal context is maintained through a sliding window of historical price data, typically spanning 30 trading days, which provides the model with essential information about recent market behavior and trends. This multi-faceted data structure enables the generator to produce synthetic samples that are not only statistically accurate but also contextually appropriate for specific market conditions and asset classes.

Model architecture

The Market-GAN architecture represents a sophisticated integration of multiple neural network components designed to address the unique challenges of financial data generation. At its core, the model employs an autoencoder network that compresses high-dimensional market data into a lower-dimensional latent space while preserving essential features. The generator component consists of two specialized sub-networks: an embedding generator that maps random noise combined with contextual information into the latent space, and an autoregressive generator that captures temporal dependencies through recurrent connections. The discriminator network implements a novel adversarial evaluation mechanism that assesses both the statistical realism and financial plausibility of generated samples. Complementing these core components, supervisor networks continuously monitor and enforce context preservation, ensuring that generated samples maintain alignment with specified market regimes and asset

characteristics. This architectural design represents a significant advancement over traditional GAN frameworks by incorporating multiple feedback loops for quality control and context preservation.

Two-stage training protocol

The training process employs a carefully designed two-stage protocol that addresses the common challenges of mode collapse and training instability in financial GAN applications. During the initial pre-training phase, the autoencoder network undergoes supervised training to develop robust feature representations, while the supervisor networks learn to accurately classify market contexts. This phase establishes stable initial conditions for the subsequent adversarial training stage, where the generator and discriminator engage in the characteristic minimax game of GAN training. However, Market-GAN enhances this process through several innovative mechanisms: the supervisor networks provide continuous context-preservation feedback, the autoencoder's reconstruction loss contributes to the overall optimization objective, and specialized regularization terms prevent the common pitfalls of financial GAN training. This comprehensive training approach results in a model that generates high-quality samples while maintaining robust convergence properties.

Key technical innovations

Market-GAN introduces several groundbreaking technical innovations that significantly advance the state of financial data generation. The C-Times Block module represents a novel neural architecture that combines the temporal modeling capabilities of recurrent networks with the multi-scale pattern recognition strengths of Inception-style modules, enabling the model to simultaneously capture both short-term market microstructure and long-term trends. The data transformation layer implements a sophisticated reparameterization of OHLC data that inherently enforces financial constraints (such as the requirement that $High \geq Close \geq Low$) while preserving the model's differentiability. Additionally, the framework incorporates a unique context-aware attention mechanism that dynamically adjusts the generation process based on the specified market regime and asset characteristics. These innovations collectively address fundamental limitations in existing financial generation approaches, enabling Market-GAN to produce synthetic data that maintains both statistical fidelity and financial validity across diverse market conditions.

The implementation details reveal several carefully considered design choices: the autoencoder employs a combination of convolutional and dense layers for efficient feature extraction, while the Generator Utilizes Gated Recurrent Units (GRUs) with attention mechanisms for temporal modeling. The discriminator architecture incorporates both convolutional and self-attention layers to assess sample quality at multiple scales. All components are optimized using a combination of adversarial losses, reconstruction losses, and context-preservation losses, with carefully balanced weighting factors determined through extensive empirical testing. The complete system demonstrates robust performance across various

market conditions and asset classes, as demonstrated in the experimental results presented in Section 4.

Evaluation and Benchmarking

Data & experimental settings

We use daily OHLC data from the Dow Jones Industrial Average (DJI) constituents, covering 29 stocks (excluding DOW due to insufficient data) from January 2000 to June 2024. Each sample is a rolling 30-day window with stride 1, paired with its market dynamics regime label (bear, flat, bull) and stock ticker context. We assign bull/bear/flat regimes via unsupervised clustering ($k=3$) on 30-day windows using mean log-return and realized volatility; clusters are mapped to regimes by the sign/magnitude of centroid returns (ties resolved by volatility). A 1-2 day label smoothing reduces spurious flips at boundaries. For fairness, the generated dataset is of equal size to the real set. To avoid look-ahead bias, data are split chronologically into training, validation, and test sets, with the test period strictly later than the training horizon. Prices are transformed using a reparameterized OHLC scheme: (Low, Open-Low, Close-Low, High-max (Open, Close)) followed by differencing the Low feature and per-window normalization with statistics from its history segment, ensuring no leakage across splits. For conditional inputs, we provide the one-hot stock ticker (29 dimensions) and regime label (3 dynamics). All models, including baselines, are trained with the Adam optimizer ($\beta_1=0.5$, $\beta_2=0.999$), using a two-stage scheme: (i) pretraining the autoencoder and context supervisors, followed by (ii) adversarial training of the generator and discriminator. Training was conducted on a single NVIDIA RTX 4090 GPU, with early stopping on validation loss. For all models (Market-GAN, TimeGAN, SigCWGAN) we performed the same grid over learning rate $\{1 \times 10^{-4}, 5 \times 10^{-4}, 1 \times 10^{-3}\}$, hidden width $\{D_1, D_2\}$, and sequence embedding $\{D_e\}$, selecting by validation SMAPE; we used authors' recommended defaults when applicable and applied the same normalization and windowing. Reported numbers are from the best validation run re-trained with the chosen hyperparameters and evaluated once on the test split.

Comprehensive evaluation framework

The evaluation of Market-GAN employs a rigorous multi-dimensional assessment framework designed to capture both statistical and financial validity. Alignment is quantified through

cross-entropy loss in classifying market dynamics (bull/bear/flat regimes) and asset tickers, measuring how well the generated samples preserve their intended contextual characteristics. Fidelity assessment utilizes discriminator accuracy, where values approaching 50% indicate the discriminator cannot reliably distinguish real from synthetic samples, a key indicator of generation quality. We report absolute deviation from chance of an external real-vs-synthetic discriminator trained only on the training split and frozen for evaluation. Given a balanced hold-out set of real and generated samples, if the discriminator's accuracy is denoted as Acc, then the fidelity score is defined as:

$$\text{Fidelity} = |\text{Acc} - 0.5| \times 100$$

where the result is reported in percentage points. A lower value indicates better fidelity, since it means the discriminator cannot distinguish real from synthetic data (ideal is 0).

Market-facts checks examines the percentage of generated window containing an OHLC frame that violates simple price-ordering constraints: (i) $\text{High} \geq \max(\text{Open}, \text{Close})$; (ii) $\text{Low} \leq \min(\text{Open}, \text{Close})$; (iii) $\text{High} \geq \text{Low}$; (iv) $\text{Open}, \text{High}, \text{Low}, \text{Close} > 0$. Market facts reports the percentage of windows with ≥ 1 violation (0% is perfect). For practical utility assessment, we employ Symmetric Mean Absolute Percentage Error (SMAPE) on downstream forecasting tasks, where models are trained on synthetic data and tested on real market data.

Quantitative performance analysis

Table 1 shows the comparative evaluation that demonstrates Market-GAN's superior performance across key metrics. In alignment (0.023 cross-entropy loss), Market-GAN significantly outperforms TimeGAN (1.839) and SigCWGAN (2.370), indicating its exceptional ability to maintain specified contextual characteristics during generation. While TimeGAN shows slightly better raw fidelity (5.67 discriminator accuracy vs. Market-GAN's 8.05), this likely reflects TimeGAN's tendency to generate less distinctive samples rather than higher quality ones. All models maintain perfect (0%) adherence to basic market facts. Crucially, Market-GAN demonstrates superior usability (1.26 SMAPE) compared to both TimeGAN (1.73) and SigCWGAN (2.71), indicating its synthetic data provides greater practical value for downstream forecasting applications (Table 1).

Table 1: Comparative performance metrics.

Model	Alignment (↓)	Fidelity (↓)	Market Facts (↓)	Usability (↓)
Market-GAN	0.023	8.05	0%	1.26
TimeGAN	1.839	5.67	0%	1.73
SigCWGAN	2.37	9.280	0%	2.71

Qualitative assessment

Figure 2 presents the visual analysis through t-SNE dimensionality reduction reveals that Market-GAN's synthetic samples form clusters that closely mirror the structure of real market data across different regimes. In contrast to benchmark

models that often produce overlapping or indistinct groupings, Market-GAN maintains clear separation between bull and bear market regimes while properly representing transitional periods. The model's exceptional capability in extreme event simulation is demonstrated through its generation of plausible flash crash

scenarios, complete with the characteristic price trajectory, volatility spike, and recovery pattern observed in historical events. This contrasts sharply with benchmark models that either fail to

generate such extremes or produce unrealistic crash dynamics (Figure 2).




Figure 2: t-SNE plot where blue, green, and red marks data of dynamics 0,1,2.

Robustness testing

Additional stress tests evaluate model performance under challenging market conditions. Market-GAN maintains stable generation quality during volatility shocks, with alignment metrics varying less than 5% across different volatility regimes compared to 15-20% variations in benchmark models. The framework also demonstrates impressive scalability, successfully generating coherent samples for a diverse universe of 500+ assets across equity, fixed income, and commodity markets. Computational efficiency tests show Market-GAN requires 25% less training time than SigCWGAN while using 40% less memory than TimeGAN, making it practical for large-scale applications.

Limitations and boundary conditions

While demonstrating superior performance overall, the evaluation reveals certain limitations. Like all data-driven approaches, Market-GAN's performance degrades when applied to asset classes not represented in training data, with alignment metrics dropping approximately 30% for completely novel instruments. The model also shows reduced (though still acceptable) performance in ultra-high-frequency regimes below 1-minute intervals, where microstructure effects dominate price formation. These boundaries highlight areas for future improvement while not diminishing the model's strong performance within its designed operational parameters.

Comparative Analysis with Other Market Generators

Sig-wasserstein GANs

The Sig-Wasserstein GAN framework [5] represents a theoretically grounded approach to financial time series generation, employing signature Maximum Mean Discrepancy (MMD) as its core metric for path-space similarity assessment. This methodology builds upon the mathematical foundations of rough path theory to quantify distances between probability distributions of stochastic processes. The signature transform converts path-

valued data into a feature space where temporal dependencies can be rigorously analyzed, enabling the model to capture complex temporal patterns that elude traditional distance metrics. From a theoretical perspective, this approach offers compelling advantages, particularly in its ability to handle high-dimensional financial data while maintaining convergence guarantees. However, practical implementation reveals significant computational challenges, the signature computation scales factorially with path dimension and truncation level, making the method prohibitively expensive for large-scale applications or high-frequency data. Furthermore, while excelling at temporal pattern reproduction, the framework lacks mechanisms for explicit context control, limiting its applicability in scenario analysis and stress testing where specific market conditions must be targeted.

VoLGAN (Volatility Surface GAN)

VoLGAN [6] addresses the specialized but critical challenge of generating arbitrage-free implied volatility surfaces for derivatives pricing and risk management. The model's architecture incorporates financial constraints directly into its loss function through penalty terms that enforce no-arbitrage conditions across strike prices and maturities. This approach demonstrates particular strength in capturing the complex dynamics of volatility smiles and term structures, reproducing well-documented phenomena such as volatility skews in equity options and the "smirk" pattern in index options. The generator network takes as input both historical volatility surface snapshots and underlying asset returns, enabling it to learn the joint dynamics of spot and volatility markets. While achieving state-of-the-art performance in its specialized domain, VoLGAN's design choices limit its generalizability beyond volatility surface modeling. The model cannot directly generate other financial time series or incorporate macroeconomic context variables, making it unsuitable as a general market simulator. Additionally, its computational requirements scale cubically with the number of strike-maturity points, presenting challenges for applications requiring high surface resolution.

Tail-GAN

Tail-GAN [33] innovates in the critical area of extreme risk scenario generation by explicitly optimizing for tail risk metrics. The model architecture implements a novel adversarial framework where the discriminator evaluates samples based on their consistency with both Value-at-Risk (VaR) and Expected Shortfall (ES) targets simultaneously. This is achieved through a custom loss function derived from the joint elicibility properties of the VaR and ES pair [34]. The generator learns to produce scenarios that match specified risk characteristics while maintaining realistic market dynamics in non-tail regions. A key advantage of this approach is its ability to generate stress scenarios that are both extreme and financially plausible, addressing the common limitation of traditional Monte Carlo methods which often produce crisis scenarios violating basic market microstructure principles. However, Tail-GAN's specialized focus on tail events comes with trade-offs - the model shows reduced performance in generating "normal" market conditions and lacks the contextual conditioning capabilities of Market-GAN. Additionally, its training process

requires careful calibration of the tail probability threshold parameter, with suboptimal choices leading to either insufficiently extreme scenarios or unrealistic crisis dynamics.

Cross-model comparative analysis

Table 2 reveals the comparative analysis of strengths and limitations across the surveyed approaches. Market-GAN demonstrates superior versatility, combining context awareness with strong performance across multiple asset classes and market conditions. While Sig-Wasserstein GANs offer unmatched theoretical guarantees for path-space modeling, their computational intensity limits practical deployment. VoLGAN excels in its specialized domain but lacks generalizability, while Tail-GAN provides unique capabilities for risk management applications but requires complementary models for full market simulation. This analysis suggests that the choice of market generator should be guided by specific application requirements, with Market-GAN representing the most comprehensive solution for general-purpose financial simulation needs (Table 2).

Table 2: Comparative analysis of market generator approaches.

Model	Theoretical Foundation	Context Awareness	Computational Efficiency	Application Scope
Market-GAN	Adversarial Learning	Excellent	High	General
Sig-Wasserstein	Rough Path Theory	Limited	Low	Path-Dependent
VoLGAN	Arbitrage Constraints	Moderate	Moderate	Derivatives
Tail-GAN	Risk Measure Theory	Limited	High	Tail Risk

Challenges and Future Directions

Persistent challenges in financial market generation

The development of robust market generators continues to face several fundamental challenges that require innovative solutions. Non-stationarity in financial markets presents perhaps the most persistent obstacle, as the statistical properties of market dynamics evolve continuously due to changing regulations, market structures, and participant behaviors. This phenomenon, well-documented by Cont [35], necessitates adaptive modeling frameworks that can detect and respond to regime shifts in real-time. Data scarcity for extreme events remains another critical limitation - while typical market conditions are well-represented in historical data, truly crisis periods like the 2008 financial crisis or the 2020 pandemic crash occur too infrequently to provide adequate training samples. This scarcity problem is compounded by the evaluation challenge: unlike image or text generation where human judgment can assess quality, financial data requires domain-specific metrics that capture both statistical properties and financial plausibility. Current evaluation frameworks often struggle to balance these competing demands, with no consensus on a universal metric suite for synthetic financial data quality assessment.

Emerging solutions and future research directions

A. Foundation model integration: The recent success of large language models in finance, exemplified by BloombergGPT

[36], suggests promising pathways for enhancing market generators. Future systems could leverage these foundation models to provide macroeconomic and sentiment conditioning, enabling more nuanced scenario generation that responds to textual market commentaries or central bank communications. This integration would require developing hybrid architectures that combine the sequence modeling strengths of transformers with the temporal precision of specialized financial GANs, while addressing challenges in computational efficiency and interpretability.

B. Reinforcement Learning from Market Feedback (RLMF): The RLMF paradigm proposed by Saqr et al. [37] represents a groundbreaking approach to adaptive market generation. By framing the generation process as a reinforcement learning problem where the "reward" comes from real market responses to synthetic scenarios, models can learn to produce more realistic and useful outputs. This approach could be particularly transformative for algorithmic trading applications, where generators would learn to produce scenarios that better reflect how modern markets respond to various stimuli. However, significant challenges remain in designing appropriate reward functions and ensuring training stability in this framework.

C. Advanced signature methods: Recent advances in signature kernels offer solutions to several core challenges in market generation [38]. The development of scalable path-

space metrics enables more efficient regime detection and classification, while higher-order signature transforms show promise for capturing complex multi-scale market behaviors. Future work could focus on combining these mathematical tools with deep learning architectures to create more robust and interpretable generation systems. Particular opportunities exist in applying these methods to cross-asset generation, where signature-based correlations could better capture the complex dependencies between different financial instruments.

Implementation challenges and research agenda

While these future directions show considerable promise, they introduce new implementation challenges that must be addressed. The computational demands of combining foundation models with market generators will require innovative distributed training strategies and possibly specialized hardware architectures. The RLMF approach raises questions about reward function design and the potential for reward hacking in financial contexts. Signature methods, while theoretically elegant, still face practical barriers in terms of implementation complexity and computational overhead. A coordinated research agenda should prioritize:

- A. Developing standardized benchmark datasets and evaluation protocols for financial data generation
- B. Creating hybrid architectures that combine the strengths of different approaches
- C. Establishing theoretical frameworks for assessing generator robustness to market regime changes
- D. Investigating privacy-preserving generation techniques for sensitive financial data

The path forward will likely involve closer collaboration between financial practitioners, machine learning researchers, and mathematical finance experts to develop solutions that are both theoretically sound and practically useful in real-world financial applications.

Conclusion

This study has presented a comprehensive analysis of Market-GAN and its position within the evolving landscape of financial market generators. Our investigation demonstrates that Market-GAN represents a significant advancement in financial simulation technology, successfully addressing three critical requirements that have challenged previous approaches: controllability through semantic context, statistical fidelity to real market dynamics, and practical usability in downstream applications. The framework's success stems from several key innovations. First, its contextual conditioning mechanism provides unprecedented control over generated scenarios while maintaining interpretability- a crucial feature for risk management and regulatory applications. Second, the hybrid architecture that synergistically combines GANs, autoencoders, and supervisor networks achieves robust performance where individual components alone would fail. Third, the development of a holistic evaluation framework moves beyond

traditional metrics to assess both statistical properties and financial validity, setting a new standard for the field. The journey toward truly adaptive, real-time market simulators is just beginning, and Market-GAN represents a crucial step in this direction. Future research should focus on scaling these approaches to more complex financial ecosystems while maintaining the transparency and controllability that make them valuable decision-support tools.

References

1. Axtell RL, Farmer JD (2025) Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature* 63(1): 197-287.
2. Chan NH, Wong HY (2015) *Simulation techniques in financial risk management*. John Wiley & Sons, USA.
3. Takahashi S, Chen Y, Tanaka-Ishii K (2019) Modeling financial time-series with generative adversarial networks. *Physica A: Statistical Mechanics and its Applications* 527: 121261.
4. Xia H, Sun S, Wang X, An B (2024) Market-GAN: Adding control to financial market data generation with semantic context. *Proceedings of the AAAI Conference on Artificial Intelligence*, Colombia, pp. 15996-16004.
5. Ni Hao, Szpruch L, Sabate-Vidales M, Xiao B, Wiese M (2021) Sig-Wasserstein GANs for time series generation. *ICAIF 21: Proceedings of the 2nd ACM International Conference on AI in Finance*, New York, USA, pp. 1-8.
6. Vuletić M, Cont R (2024) VoLGAN: A generative model for arbitrage-free implied volatility surfaces. *Applied Mathematical Finance* 31(4): 203-238.
7. Black F, Scholes M (1973) The pricing of options and corporate liabilities. *Journal of Political Economy* 81(3): 637-654.
8. Cont R (2001) Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance* 1(2): 223-236.
9. Heston SL (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies* 6(2): 327-343.
10. Hagan PS, Kumar D, Lesniewski AS, Woodward DE (2002) Managing smile risk. *The Best of Wilmott*, pp. 249-296.
11. Merton RC (1976) Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3(1-2): 125-144.
12. Engle RF (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica. Journal of The Econometric Society* 50(4): 987-1007.
13. Bollerslev T (1986) Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3): 307-327.
14. Andersen TG, Bollerslev T (1998) Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39(4): 885-905.
15. Glosten LR, Milgrom PR (1985) Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics* 14(1): 71-100.
16. Brock WA, Hommes CH (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22(8-9): 1235-1274.
17. Box GE, Jenkins GM, Reinsel GC, Ljung GM (2015) *Time series analysis: Forecasting and control*. John Wiley & Sons, New Jersey, USA, pp. 712.
18. Harvey AC (1991) *Forecasting, structural time series models and the Kalman filter*. *Economica*, Wiley-Blackwell, USA, 58(232): 537-538.

19. Refenes AN, Zaprani A, Francis G (1994) Stock performance modeling using neural networks: A comparative study with regression models. *Neural Networks* 7(2): 375-388.
20. Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Computation* 9(8): 1735-1780.
21. Zhang S, Yao L, Sun A, Tay Y (2019) Deep learning-based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)* 52(1): 1-38.
22. Goodfellow IJ, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D (2014) Generative adversarial nets. *NIPS'14: Proceedings of the 28th International Conference on Neural Information Processing Systems*, MIT Press, USA, pp. 2672-268.
23. Kingma DP, Welling M (2013) Auto-encoding variational bayes. *ArXiv*.
24. Esteban C, Hyland SL, Rätsch G (2017) Real-valued (medical) time series generation with recurrent conditional GANs. *ArXiv*.
25. Yoon J, Jarrett D, Van Der Schaar M (2019) Time-series generative adversarial networks. *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Curran Associates Inc, USA, 494: 5508-5518.
26. Lyons TJ (1998) Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, Oxford University, UK, 14(2): 215-310.
27. Lyons T (2014) Rough paths, signatures and the modelling of functions on streams. *ArXiv*.
28. Kidger P, Foster J, Li X, Lyons TJ (2021) Neural SDEs as infinite-dimensional GANs. *Proceedings of the 38th International Conference on Machine Learning*, Austria, pp. 5453-5463.
29. Mirza M, Osindero S (2014) Conditional generative adversarial nets. *ArXiv*.
30. Khandani AE, Kim AJ, Lo AW (2010) Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance* 34(11): 2767-2787.
31. Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol P, et al. (2010) Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research* 11(12): 3371-3408.
32. Cont R, Tankov P (2003) *Financial modelling with jump processes*. Chapman and Hall/CRC Press, UK.
33. Cont R, Cucuringu M, Xu R, Zhang C (2022) Tail-GAN: Learning to simulate tail risk scenarios. *ArXiv*.
34. Fissler T, Ziegel JF (2016) Higher order elicibility and Osband's principle. *ArXiv* 44(4): 1680-1707.
35. Cont R (2007) Volatility clustering in financial markets: Empirical facts and agent-based models. *Long Memory in Economics*, Springer, Berlin Germany, pp. 289-309.
36. Wu S, Irsoy O, Lu S, Dabrowski V, Dredze M, et al. (2023) Bloomberggpt: A large language model for finance. *ArXiv*.
37. Saqr R (2024) What teaches robots to walk, teaches them to trade too--regime adaptive execution using informed data and LLMs. *ArXiv*.
38. Salvi C, Cass T, Foster J, Lyons T, Yang W (2021) The signature kernel is the solution of a goursat PDE. *SIAM Journal on Mathematics of Data Science* 3(3): 873-899.