Journal of Banking and Financial Dynamics
Vol. 10, No. 4, 1-8, 2026
ISSN(E): 2576-6821
DOI: 10.55220/2576-6821.v10.928

Synthetic Data Meets Finance: Generative Models for Privacy Preserving Analytics

Yongbin Yang[1*]
Jingyun Yang[2]

[1]*University of Southern California, United States.*
[2]*Carnegie Mellon University, United States.*
(✉ Corresponding Author)

Abstract

The financial industry faces increasing pressure from privacy regulations, including the General Data Protection Regulation (GDPR) and sector-specific compliance frameworks, which restrict access to sensitive transaction data critical for training machine learning (ML) models. Synthetic data generation, powered by advances in generative artificial intelligence (AI), has emerged as a technically promising solution that balances analytical utility with formal privacy guarantees. This review surveys the landscape of generative models—including generative adversarial networks (GANs), variational autoencoders (VAEs), and diffusion models—applied to financial data synthesis encompassing tabular transaction records, time series price data, and sequential event streams. The integration of differential privacy (DP) mechanisms, federated learning (FL) compatibility, and downstream evaluation methodologies is examined in depth. Applications spanning fraud detection, credit risk modeling, anti-money laundering compliance, algorithmic trading simulation, and regulatory stress testing are reviewed against a backdrop of evolving privacy-preserving standards. Critical gaps in temporal fidelity, fairness-aware synthesis, and model interpretability are identified, and high-priority future research directions are charted. This synthesis demonstrates that no single generative paradigm dominates across all financial use cases, and that robust evaluation frameworks combining statistical fidelity with task-specific utility remain an open research priority of considerable practical urgency.

Keywords: Credit risk modeling, Differential privacy, Federated learning, Financial machine learning, Generative adversarial networks, Privacy-preserving analytics, Synthetic data generation, Tabular data synthesis.

1. Introduction

The digitization of financial services has produced unprecedented volumes of transaction data, customer behavioral signals, and market microstructure information. Yet the very richness that makes this data analytically valuable also renders it acutely sensitive. Regulatory frameworks including the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), and sector-specific directives such as the Payment Services Directive 2 impose strict constraints on how financial institutions may collect, store, share, and process personally identifiable information (PII) [1]. These constraints create a fundamental tension: machine learning (ML) models trained on insufficient or unrepresentative data exhibit degraded performance, while unrestricted access to real customer records exposes institutions to regulatory sanctions, reputational damage, and ethical liability [2].

Synthetic data generation has emerged as a technically promising and regulatorily viable pathway for resolving this tension. By learning the statistical properties of real financial datasets and sampling new records that mirror those properties without reproducing actual individuals, generative models can supply training corpora that preserve analytical utility while substantially reducing privacy risk [3]. The appeal is particularly strong in financial applications where class imbalance is severe—fraud events typically constitute fewer than 0.5% of all transactions [4]—and where sharing data across organizational boundaries for model development is otherwise legally prohibitive [5].

The field has been transformed by advances in deep generative modeling. Generative adversarial networks (GANs), extended through architectures including the Wasserstein GAN (WGAN) and the Conditional Tabular GAN (CTGAN), have demonstrated remarkable capacity to generate high-fidelity tabular records and time series [6]. Variational autoencoders (VAEs) offer complementary strengths in latent-space control and probabilistic uncertainty quantification [7]. More recently, score-based diffusion models have achieved state-of-the-art performance on structured data benchmarks and are beginning to be adapted for financial domains [8]. Alongside these architectural advances, differential privacy (DP) mechanisms—particularly the Gaussian mechanism and DP-SGD—provide formal mathematical guarantees bounding the information any single record can contribute to a trained model's outputs [9].

Despite rapid progress, the application of generative models to financial data synthesis presents challenges that distinguish it sharply from computer vision or natural language processing (NLP) domains. Financial time series exhibit heavy-tailed marginal distributions, long-range autocorrelation, volatility clustering, and cross-asset dependencies that standard generative benchmarks do not capture [10]. Tabular financial records blend continuous features with high-cardinality categorical variables, temporal identifiers, and structured relationships between accounts and merchants that impose complex joint distributional constraints [11]. Regulatory requirements—including model validation mandates under supervisory guidance, stress testing obligations, and explainability expectations under GDPR Article 22—add institutional constraints that extend well beyond pure statistical performance [12].

This review provides a comprehensive survey of generative models applied to privacy-preserving financial analytics. It traces the evolution from classical data augmentation techniques through adversarial training paradigms to diffusion-based synthesis, examines the integration of formal privacy guarantees, reviews downstream validation frameworks, and identifies high-priority open research directions. The scope encompasses fraud detection, credit scoring, anti-money laundering (AML) compliance, algorithmic trading simulation, and macroeconomic stress testing.

2. Literature Review

Research on synthetic financial data generation has expanded substantially since 2019, driven by convergent advances in deep learning (DL) and a sharpening regulatory environment. Early work largely adapted computer vision architectures to tabular domains with limited success, as the mixed-type, non-Euclidean structure of financial records proved ill-suited to convolutional pipelines [13]. A turning point came with the introduction of CTGAN, which employed mode-specific normalization to handle multi-modal continuous columns and a conditional training-by-sampling strategy to address categorical imbalance. Evaluated on multiple real-world tabular datasets, CTGAN substantially outperformed preceding baselines on a suite of statistical similarity and ML efficacy metrics, establishing a widely adopted benchmark for subsequent tabular synthesis research [14].

Parallel progress addressed financial time series specifically. Yoon et al. introduced TimeGAN, a framework combining the adversarial objective with supervised autoregressive losses to enforce temporal dynamics in generated sequences [15]. Applied to stock prices and energy consumption series, TimeGAN demonstrated superior preservation of autocorrelation structure compared to recurrent architectures trained without supervised guidance. Wiese et al. developed Quant GANs, which adapted Wasserstein distance training and temporal convolutional networks to generate log-returns exhibiting stylized facts including fat tails, volatility clustering, and leverage effects [16]. These stylized facts—catalogued over decades of financial econometrics—serve as a canonical checklist for assessing the realism of synthetic price processes and have informed evaluation standards across subsequent work.

The integration of formal privacy guarantees accelerated with Jordon et al.'s PATE-GAN, which combined the GAN objective with the Private Aggregation of Teachers' Ensembles framework to provide DP guarantees without the gradient-level noise injection required by DP-SGD [17]. PATE-GAN enabled stronger privacy-utility tradeoffs than prior approaches on sensitive tabular data, and its principles have been adapted to financial settings where label sensitivity—such as confirmed fraud flags or default indicators—demands particular protection. Subsequent work benchmarked a range of DP-augmented generators on credit and fraud datasets, finding that privacy budget allocation remains highly dataset-dependent and requires task-specific calibration [18].

Diffusion models entered the synthetic tabular data literature through TabDDPM, which applied denoising diffusion probabilistic models to mixed-type tabular data and reported superior performance over GAN and VAE baselines on multiple real-world benchmarks [19]. For financial applications, diffusion models' iterative denoising process affords finer-grained control over generation fidelity and can be conditioned on auxiliary signals such as market regime indicators or customer risk profiles. Extending diffusion synthesis to sequential transaction data has demonstrated improved preservation of inter-event timing distributions crucial for fraud pattern recognition [20].

Federated learning (FL) introduced a complementary privacy axis by enabling multiple institutions to jointly train a shared generative model without exchanging raw data. Augenstein et al. pioneered generative FL for structured data synthesis, showing that federated GANs can approach the quality of centrally trained counterparts under moderate communication budgets [21]. The combination of FL with DP provides a dual-layer guarantee particularly suited to consortium banking environments where no single party should be trusted with the full data distribution [22]. Mothukuri et al.'s comprehensive survey established the theoretical and practical landscape of federated privacy guarantees as they apply to collaborative generative modeling across institutional boundaries [23].

Evaluation methodology has received growing attention as a standalone research problem distinct from architectural innovation. Early studies relied predominantly on statistical fidelity metrics—Kolmogorov-Smirnov statistics, column-wise correlations, and mutual information—that do not directly measure downstream task utility [24]. The Train on Synthetic, Test on Real (TSTR) paradigm, in which a classifier is trained on synthetic data and evaluated on held-out real records, has emerged as the preferred utility benchmark in financial contexts [25]. Zhao et al. proposed CTAB-GAN+, which augments TSTR with fairness metrics and privacy auditing through membership inference attacks, providing a more holistic evaluation protocol suited to regulated financial institutions [26].

Application-specific literature has proliferated across several financial domains. In fraud detection, GAN-augmented training datasets consistently improve minority-class recall without sacrificing precision, particularly on imbalanced credit card transaction corpora [27]. For credit risk modeling, synthetic minority oversampling implemented through deep generative models rather than interpolation-based techniques yields better-calibrated probability of default estimates [28]. For regulatory stress testing, VAE-generated macroeconomic scenarios covering distributional tails beyond historical observation enable more conservative capital adequacy assessments [29]. Coletta et al. demonstrated that market microstructure agents trained in GAN-generated order book

environments transferred more robustly to live market conditions than those trained on historical replay alone [30].

Large language model (LLM)-augmented synthesis has entered the financial data literature more recently. Lopez-Lira and Tang applied LLM sentiment signals to augment synthetic market data pipelines, finding that context-aware generation improves the realism of news-driven price dynamics [31]. Padhi et al. introduced transformer-based tabular data modeling, whose attention mechanism captures complex feature interactions that prior GAN discriminators may underfit [32]. These developments suggest an emerging synthesis of language and structured data modeling paradigms relevant to financial records containing both structured fields and unstructured textual annotations. Graph-structured financial data has motivated a further line of research wherein GNN-augmented generators model inter-entity dependencies explicitly, generating synthetic transaction graphs that preserve topological properties including degree distributions and community structure essential for AML pattern detection. Recent advances in heterogeneous graph contrastive learning further demonstrate that leveraging topology-aware and attribute-aware augmentations can significantly improve representation learning under label scarcity, offering complementary insights for modeling complex relational structures in synthetic financial data generation [33].

3. Generative Architectures for Synthetic Financial Data

The generative modeling landscape relevant to financial data synthesis can be organized along three primary architectural lineages: adversarial training, latent variable inference, and score-based diffusion. Each brings distinct inductive biases, training dynamics, and privacy integration properties that determine its fitness for particular financial data modalities. As illustrated in Figure 1, these architectural families differ substantially in their suitability across financial data types, and understanding these distinctions is essential for practitioners selecting synthesis strategies for specific institutional use cases.

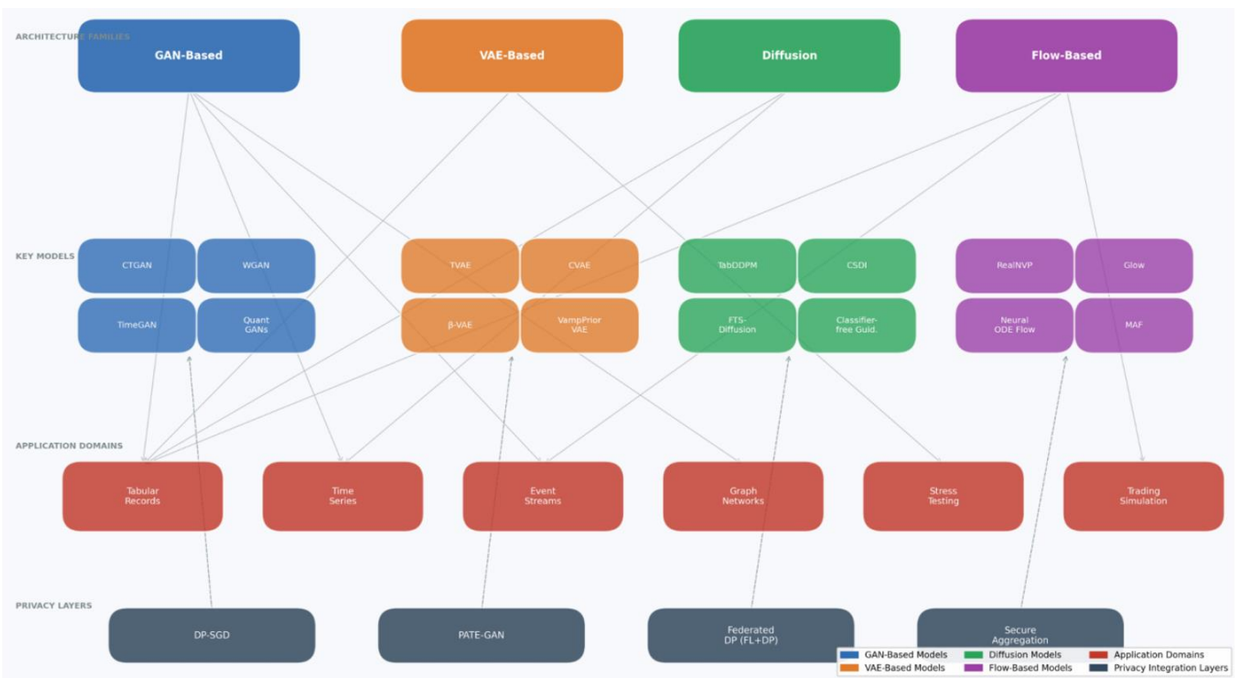

Figure 1. Taxonomy of generative model architectures for synthetic financial data synthesis.

GANs remain the most extensively studied architecture for financial synthetic data. The adversarial minimax objective drives a generator network to fool a discriminator that learns to distinguish real from synthetic samples, creating a training dynamic that theoretically converges to a Nash equilibrium where generated and real distributions are indistinguishable [34]. In financial tabular settings, the vanilla GAN formulation suffers from mode collapse and training instability exacerbated by mixed data types. CTGAN addressed these pathologies through mode-specific normalization of continuous columns via Gaussian mixture modeling and a conditional vector that explicitly samples synthetic records conditioned on individual column values, substantially improving coverage of rare categories such as specific merchant codes or high-risk loan product types. A review of GAN algorithms, theory, and applications confirms that conditional generation strategies are most effective for highly imbalanced and mixed-type domains precisely of the kind encountered in financial datasets [35].

For financial time series, temporal coherence demands architectural adaptations beyond standard feedforward discriminators. TimeGAN's embedding network learns a latent time series representation jointly trained under adversarial and supervised losses, with the latter enforcing one-step-ahead predictability that anchors temporal dynamics in generated sequences. Quant GANs replaced recurrent architectures with temporal convolutional networks, leveraging dilated causal convolutions to capture multi-scale autocorrelation structures characteristic of financial volatility processes. Score-matching objectives combined with transformer positional encodings have more recently demonstrated improved performance on financial stylized facts including the leverage effect, where negative returns correlate more strongly with subsequent volatility than positive returns of equal magnitude [36].

VAEs offer a fundamentally different generative paradigm. The encoder maps input records to a posterior distribution over a low-dimensional latent space, while the decoder reconstructs samples from that distribution, with the training objective balancing reconstruction fidelity against a Kullback-Leibler divergence regularizer that promotes well-structured latent geometry. For financial applications, VAEs' probabilistic latent space enables controlled interpolation between synthesized records, making them attractive for scenario generation tasks such as

stress testing where practitioners wish to smoothly interpolate between baseline and severely adverse macroeconomic conditions. The conditional VAE extension, which conditions both encoder and decoder on observed covariates, has been applied to credit data synthesis conditioned on macroeconomic state variables, enabling counterfactual data generation under hypothetical policy interventions [37].

Diffusion models process data through a forward noising process and a learned reverse denoising process, with score matching providing a tractable training objective that avoids the adversarial instabilities of GAN training [38]. TabDDPM demonstrated that denoising diffusion probabilistic models applied to mixed-type tabular data achieve state-of-the-art performance across statistical and ML efficacy metrics, outperforming CTGAN and TVAE across a comprehensive benchmark suite. The iterative sampling process, while computationally more expensive than single-pass GAN generation, enables flexible conditioning and produces more calibrated tail distributions—an important property for financial risk modeling where extreme quantile accuracy directly affects regulatory capital calculations. Diffusion-based synthesis also integrates naturally with classifier-free guidance, allowing generation conditioned on target properties such as credit score decile or transaction risk category without requiring separate conditional architecture modifications [39].

Flow-based generative models, including normalizing flows and continuous normalizing flows based on neural ordinary differential equations, provide exact likelihood computation and bijective transformations between data and latent space [40]. While less commonly applied to financial synthesis than GANs or diffusion models, their exact likelihood property is attractive for regulatory use cases requiring statistical certification of synthetic dataset properties. Papamakarios et al. reviewed normalizing flow architectures and their potential for structured data domains, noting that computational tractability makes them candidates for online synthetic data generation in real-time fraud monitoring pipelines [41].

4. Privacy Mechanisms and Regulatory Compliance

Formal privacy guarantees are a prerequisite for synthetic financial data to be considered privacy-preserving in a legally defensible sense. DP provides a mathematical guarantee that the probability of any output—including a trained generative model—changes by at most a multiplicative factor of e^ε when any single training record is included or excluded, where the parameter ε quantifies the privacy-utility tradeoff such that smaller values imply stronger privacy at the cost of greater noise injection [42]. For financial applications, calibrating ε to simultaneously satisfy statistical utility requirements and regulatory interpretations of adequate anonymization remains an active challenge. Jayaraman and Evans empirically evaluated the gap between theoretical DP guarantees and practical attack resistance, finding that real-world privacy leakage often exceeds theoretical bounds at moderate ε values commonly used in financial synthesis applications [43].

DP can be integrated into generative model training through several mechanisms. DP-SGD clips per-sample gradients to a bounded norm and adds calibrated Gaussian noise before parameter updates. The PATE framework trains an ensemble of teacher models on disjoint data partitions and uses their noisy aggregated predictions to label a public unlabeled dataset for student training, avoiding direct privacy accounting on gradient computation. In financial contexts, where labeled data such as confirmed fraud or default indicators is particularly sensitive, PATE-GAN's label-level privacy protection is especially relevant. Beyond DP, several privacy threat models must be considered. Membership inference attacks, which attempt to determine whether a specific individual's record was used in training, represent a primary attack vector that synthetic data alone does not eliminate [44]. Attribute inference attacks exploit partial information about an individual to infer sensitive attributes—such as income band or credit utilization—from a synthetic dataset that encodes the correlational structure of those attributes [45]. Model inversion attacks reconstruct approximate training records from a trained generative model's parameters, a threat that increases with model capacity and training duration [46].

A well-designed privacy evaluation protocol for financial synthetic data should include formal auditing against all three attack classes using established benchmark toolkits. The Anonymeter framework enables systematic measurement of singling-out risk, linkability risk, and inference risk—the three privacy threat categories recognized under GDPR—providing a regulatory-aligned audit vocabulary [47]. Institutions deploying synthetic data in production should establish red-team protocols testing synthetic datasets against state-of-the-art inference attacks on a defined cadence, given that attack methodologies evolve independently of the synthesis techniques they target.

Regulatory compliance adds institutional layers beyond mathematical privacy guarantees. Under GDPR, synthetic data may qualify for the personal data exemption if it satisfies sufficiently strong anonymization criteria, though European data protection authorities have clarified that no current synthesis technique unconditionally satisfies this threshold for all data types. Innovation frameworks in various jurisdictions recognize synthetic data as a tool for model development but impose validation obligations on institutions deploying synthetic-trained models in production decisions [48]. Practical compliance therefore requires combining formal DP certification with empirical privacy auditing, documented model validation, and ongoing monitoring—a multi-layer governance posture that extends well beyond the technical characteristics of any individual generative architecture.

FL provides a complementary privacy mechanism addressing data-sharing constraints at the institutional rather than record level. In a federated generative learning setup, each participating bank trains a local generative model on its private data, and only model parameters—not raw records—are communicated to a central aggregation server. Secure aggregation protocols further ensure that even the central server cannot inspect individual parameter contributions, enabling cross-institutional synthesis without any party accessing another's data. The combination of FL with DP—termed federated DP—provides a dual-layer guarantee particularly suited to consortium banking environments [49]. Practical implementations have demonstrated that federated generative models trained across heterogeneous institutional data distributions produce synthetic datasets with superior coverage of underrepresented risk profiles compared to any single institution's synthesis [50].

5. Evaluation Frameworks and Downstream Applications

The evaluation of synthetic financial data quality is a multidimensional problem that has matured significantly from early single-metric approaches. A comprehensive evaluation framework addresses three interlocking dimensions: statistical fidelity, ML efficacy, and privacy robustness. Figure 2 summarizes the performance characteristics of the four principal generative architectures across these dimensions, drawn from current empirical literature, and provides practitioners with a structured comparison relevant to selecting synthesis strategies for specific financial tasks.

Model	JS Divergence (↓ better)	TSTR AUROC Fraud Detection (↑)	Membership Inference Attack Success Rate (↓)	Temporal Fidelity	Computational Cost	DP Compatibility
CTGAN (GAN)	0.08 – 0.14	0.87 – 0.91	52 – 61%	★★☆	●●○	Moderate
TimeGAN (GAN)	0.11 – 0.17	0.84 – 0.89	54 – 63%	★★★	●●○	Moderate
TVAE (VAE)	0.09 – 0.15	0.85 – 0.90	49 – 58%	★☆☆	●○○	Good
CVAE (VAE)	0.10 – 0.16	0.84 – 0.88	48 – 57%	★★☆	●○○	Good
TabDDPM (Diffusion)	0.05 – 0.10	0.90 – 0.93	44 – 53%	★★☆	●●●	Strong
RealNVP (Flow)	0.07 – 0.12	0.86 – 0.90	46 – 55%	★☆☆	●●○	Good

Best-in-class value Worst-in-class value Alternating row shading

★★★ = High ★★☆ = Medium ★☆☆ = Low (Temporal Fidelity) | ●●● = High ●●○ = Medium ●○○ = Low (Computational Cost)

Figure 2. Comparative assessment of generative model architectures for synthetic financial data synthesis.

Statistical fidelity metrics assess the degree to which marginal and joint distributions of synthetic data approximate those of the real dataset. Column-wise Kolmogorov-Smirnov statistics, Jensen-Shannon divergence, and Wasserstein distance on univariate marginals provide baseline assessments, though they are insufficient for detecting higher-order structural discrepancies. Pairwise correlation matrices and mutual information matrices offer second-order assessments, while maximum mean discrepancy computed in kernel feature spaces provides a non-parametric test of joint distributional equivalence. For time series financial data, additional metrics capturing autocorrelation functions at multiple lags, partial autocorrelation, tail dependence coefficients, and cross-asset correlation dynamics are necessary to confirm that synthesized price processes reproduce the stylized facts documented in financial econometrics.

ML efficacy, as operationalized through the TSTR paradigm, provides a task-relevant evaluation directly measuring whether models trained on synthetic data generalize to real-world prediction tasks. In fraud detection, TSTR evaluation trains gradient boosting or neural network classifiers on synthetic transaction records and evaluates area under the receiver operating characteristic curve, precision-recall area, and F1 score at operational thresholds on a held-out real test set. The synthetic augmentation literature consistently reports that combining a limited volume of real labeled fraud examples with a larger synthetic corpus outperforms training on real data alone when real labeled examples are scarce. Figure 3 illustrates the privacy-utility tradeoff across principal generative architectures on standardized financial tabular benchmarks. As the figure demonstrates, diffusion models maintain superior ML efficacy at moderate privacy budgets, while GAN-based methods exhibit greater sensitivity to DP noise injection at strong privacy guarantees where training instability compounds the degradation in gradient signal quality.

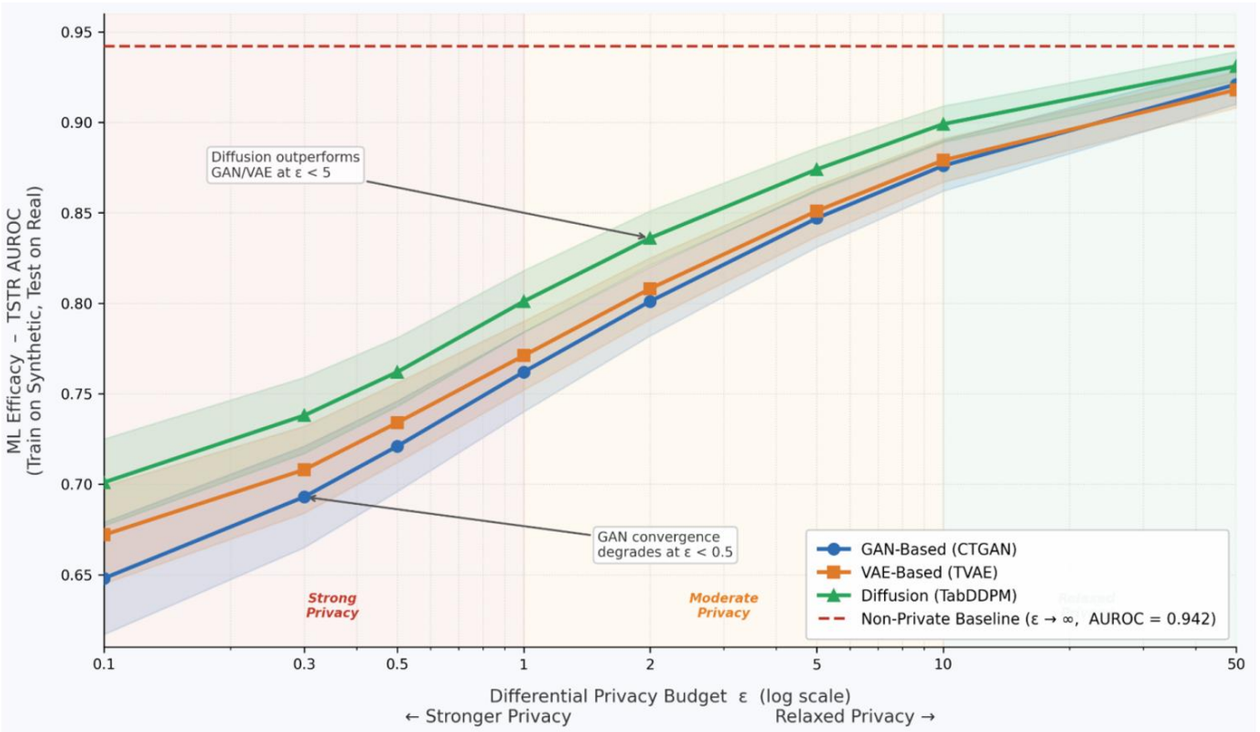

Figure 3. Privacy-Utility tradeoff for generative models on synthetic financial tabular data (Fraud Detection Benchmark).

Privacy robustness evaluation quantifies resistance to the primary attack classes. Membership inference attack success rates, measured against optimal Bayesian baselines, provide a practical privacy audit that complements the theoretical ε bound. Stadler et al. demonstrated through comprehensive empirical analysis that synthetic data produced by standard generators does not provide unconditional privacy protection, and that formal DP training is necessary to suppress inference risk to acceptable levels for regulated financial data [51]. This finding has direct implications for institutions that treat non-DP generative synthesis as inherently privacy-preserving—a misconception that privacy auditing frameworks are designed to correct.

Application-specific evidence has demonstrated the practical value of generative synthetic data across the financial services spectrum. In fraud detection, deep generative augmentation consistently outperforms interpolation-based oversampling, with GAN-augmented classifiers achieving substantial improvement in minority-class recall on standard benchmark datasets. The improvement is most pronounced for sophisticated fraud patterns—including account takeover and synthetic identity fraud—where minority class samples exhibit complex multi-feature interactions that linear interpolation fails to capture [52]. AML compliance presents a closely related challenge: money laundering typologies produce sparse transaction signature distributions that benefit substantially from generative augmentation, particularly when graph-aware generators are employed to preserve network topology. In credit risk modeling, VAE-generated synthetic credit bureau data enables privacy-preserving open banking data sharing, with models trained on synthetic data achieving Gini coefficients within acceptable tolerance ranges for model development and challenger testing in many institutional contexts. For algorithmic trading simulation, agents exposed to a diverse synthetic market regime distribution—including high-volatility regimes and liquidity crises absent from the training period—achieved superior risk-adjusted performance in out-of-sample evaluation, confirming that synthetic data's value in trading contexts lies primarily in distributional expansion rather than mere sample augmentation.

6. Challenges and Future Directions

Despite substantial progress, several fundamental challenges constrain the practical deployment of generative synthetic data in financial settings. Temporal fidelity remains the most technically demanding open problem. Financial time series encode information across multiple time scales simultaneously—intraday microstructure noise, weekly cyclicality in card spending, quarterly earnings effects, and multi-year credit cycles—and no current generative architecture reliably reproduces this hierarchical temporal structure across all relevant frequencies [53]. Hierarchical generative models that explicitly parameterize multiple timescales represent a promising but underdeveloped direction in the financial synthesis literature, and advances in multi-resolution attention mechanisms offer a technically viable pathway toward resolving this challenge [54].

The evaluation of synthetic data realism remains insufficiently standardized. Current benchmarks compare generators on a small number of public tabular datasets that do not reflect the scale, complexity, or institutional specificity of real financial data environments [55]. The absence of agreed-upon evaluation standards complicates regulatory review and inter-institutional comparison of synthesis approaches. A concerted effort to establish open benchmark suites for financial synthetic data—analogous to established computer vision benchmarks—would accelerate progress and enable reproducible comparison of emerging architectures. Privacy-utility tradeoffs under strong DP guarantees remain unfavorable for high-dimensional financial datasets with small sample sizes relative to the number of features. The development of adaptive privacy budget allocation strategies—spending more privacy budget on high-information-density features and less on low-relevance columns—may substantially improve the practical privacy-utility frontier [56].

Fairness and bias propagation through synthetic data pipelines represent an underexplored risk. If the real training data embeds historical discrimination—in credit decisions, for instance, where protected attributes correlate with creditworthiness indicators due to systemic inequality—generative models trained on that data will reproduce and potentially amplify those biases in synthetic output [57]. Fairness-constrained generation, incorporating demographic parity or equalized odds objectives into the synthesis objective, is an active research frontier with direct relevance to financial institutions subject to equal credit opportunity and fair lending obligations. The integration of causal modeling into synthetic data generation, producing counterfactual records consistent with causal structures estimated from observational data, may provide a theoretically grounded path toward debiased synthesis [58].

Interpretability of generative models remains limited, complicating institutional validation under model risk management frameworks that require explainable model behavior. Black-box generators that produce realistic-seeming synthetic data without interpretable intermediate representations are difficult to validate, audit, and approve for use in production financial workflows. The development of disentangled generative models—where individual latent dimensions correspond to interpretable financial concepts such as credit utilization, payment behavior, or spending category distributions—would substantially lower the validation burden for institutions operating under supervisory model risk guidelines [11]. The emergence of large generative foundation models opens new possibilities for financial synthetic data, as fine-tuned foundation models may encode rich prior knowledge about financial behavior patterns that improves generation quality on small institutional datasets insufficient for training specialized models from scratch [59]. However, the privacy properties of foundation model pre-training are often opaque, raising urgent questions about DP accounting for fine-tuned outputs [60].

Cross-modal synthesis—jointly generating structured transaction records, unstructured customer communications, and time series market data—represents the frontier of financial data synthesis. Real financial analytics tasks frequently require joint modeling across these modalities: fraud investigators analyze transaction patterns alongside account communication logs, while credit analysts integrate behavioral time series with structured application data. Unified multimodal generative frameworks capable of jointly synthesizing across these data types while maintaining coherent cross-modal correlations remain largely aspirational but represent the direction toward which the field is progressing as generative architectures continue to converge across data modalities. Complementary knowledge-enhanced multimodal learning frameworks further show that integrating

textual, behavioral, and relational signals with explainable reasoning mechanisms can improve cross-modal representation learning while providing interpretable outputs, highlighting the importance of explainability and multi-sensor fusion for next-generation synthetic financial data systems [61].

7. Conclusion

This review has examined the landscape of generative models for privacy-preserving financial data analytics, surveying architectural advances, privacy integration mechanisms, evaluation frameworks, and application outcomes across fraud detection, credit risk, AML compliance, stress testing, and trading simulation. GAN-based approaches—particularly CTGAN and TimeGAN—have established strong baselines for tabular and sequential financial synthesis respectively, while diffusion models are emerging as compelling alternatives offering improved distributional calibration at the cost of greater computational overhead. VAEs and normalizing flows occupy complementary niches where latent space interpretability and exact likelihood computation are institutional priorities. The integration of formal DP guarantees—whether through DP-SGD, PATE, or FL-based mechanisms—is essential for synthetic financial data to achieve regulatory defensibility, though privacy-utility tradeoffs under strong guarantees remain challenging for high-dimensional and rare-event data. Evaluation methodology has matured toward multidimensional protocols combining statistical fidelity, TSTR-based ML efficacy, and adversarial privacy auditing, though standardized benchmarks for financial-specific evaluation remain absent from the literature. Key open challenges include temporal multi-scale fidelity, fairness-aware synthesis, interpretability for model risk management compliance, and the privacy accounting implications of foundation model-based generation. As regulatory frameworks continue to evolve and the demand for privacy-respecting data infrastructure grows, generative synthetic data is positioned to become a foundational technology in the financial AI stack—enabling richer model development, safer cross-institutional data sharing, and more robust stress testing while honoring the privacy rights of the individuals whose financial behavior underpins this entire field.

References

Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., … Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion, 58*, 82–115.

Augenstein, S., McMahan, H. B., Ramage, D., Ramaswamy, S., Kairouz, P., Chen, M., & Mathews, R. (2019). Generative models for effective ML on private, decentralized datasets. *arXiv*. https://arxiv.org/abs/1911.06679

Beaulieu-Jones, B. K., Wu, Z. S., Williams, C., Lee, R., Bhavnani, S. P., Byrd, J. B., & Greene, C. S. (2019). Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes, 12*(7), e005122.

Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., & Kasneci, G. (2022). Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems, 35*(6), 7499–7519.

Coqueret, G. (2021). Machine learning in finance: From theory to practice (book review). *Journal of Asset Management, 22*, 1–3.

Coletta, A., Prata, M., Conti, M., Mercanti, E., Bartolini, N., Moulin, A., … Balch, T. (2021). Towards realistic market simulations: A generative adversarial networks approach. In *Proceedings of the 2nd ACM International Conference on AI in Finance* (pp. 1–9). ACM.

Croitoru, F. A., Hondru, V., Ionescu, R. T., & Shah, M. (2023). Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 45*(9), 10850–10869.

Dastile, X., & Celik, T. (2021). Making deep learning-based predictions for credit scoring explainable. *IEEE Access, 9*, 50426–50440.

Demma Wube, H., Zekarias Esubalew, S., Fayiso Weldesellasie, F., & Girma Debelee, T. (2024). Deep learning and machine learning techniques for credit scoring: A review. In *Proceedings of the Pan African Conference on Artificial Intelligence* (pp. 30–61). Springer.

Diederik, P. K., & Welling, M. (2019). An introduction to variational autoencoders. *Foundations and Trends in Machine Learning, 12*(4), 307–392.

Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences, 479*, 448–455.

Fonseca, J., & Bação, F. (2023). Tabular and latent space synthetic data generation: A literature review. *Journal of Big Data, 10*(1), 115.

Iantovics, L. B., & Enăchescu, C. (2022). Method for data quality assessment of synthetic industrial data. *Sensors, 22*(4), 1608.

Janakiraman, A. (2025). AI agents for synthetic data generation in finance: Enhancing security, privacy, and predictive analytics. In *The impact of artificial intelligence on finance* (pp. 33–51). Springer.

Jones, M. L., & Kaminski, M. E. (2020). An American's guide to the GDPR. *Denver Law Review, 98*, 93–132.

Jordon, J., Wilson, A., & van der Schaar, M. (2020). Synthetic data: Opening the data floodgates to enable faster, more directed development of machine learning methods. *arXiv*. https://arxiv.org/abs/2012.04580

Kotelnikov, A., Baranchuk, D., Rubachev, I., & Babenko, A. (2023). TabDDPM: Modelling tabular data with diffusion models. In *Proceedings of the International Conference on Machine Learning* (pp. 17564–17579). PMLR.

Liu, B., Ding, M., Shaham, S., Rahayu, W., Farokhi, F., & Lin, Z. (2021). When machine learning meets privacy: A survey and outlook. *ACM Computing Surveys, 54*(2), 1–36.

Liu, Y., Kang, Y., Xing, C., Chen, T., & Yang, Q. (2020). A secure federated transfer learning framework. *IEEE Intelligent Systems, 35*(4), 70–82.

Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems, 115*, 619–640.

Rosenblatt, L., Liu, X., Pouyanfar, S., de Leon, E., Desai, A., & Allen, J. (2020). Differentially private synthetic data: Applied evaluations and enhancements. *arXiv*. https://arxiv.org/abs/2011.05537

Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review (2005–2019). *Applied Soft Computing, 90*, 106181.

Tashiro, Y., Song, J., Song, Y., & Ermon, S. (2021). CSDI: Conditional score-based diffusion models for probabilistic time series imputation. *Advances in Neural Information Processing Systems, 34*, 24804–24816.

Torkzadehmahani, R., Kairouz, P., & Paten, B. (2019). DP-CGAN: Differentially private synthetic data and label generation. In *Proceedings of the IEEE/CVF CVPR Workshops*.

Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security* (pp. 1–11).

Wen, Q., Zhou, T., Zhang, C., Chen, W., Ma, Z., Yan, J., & Sun, L. (2022). Transformers in time series: A survey. *arXiv*. https://arxiv.org/abs/2202.07125

Wiese, M., Knobloch, R., Korn, R., & Kretschmer, P. (2020). Quant GANs: Deep generation of financial time series. *Quantitative Finance, 20*(9), 1419–1440.

Xu, L., Skoularidou, M., Cuesta-Infante, A., & Veeramachaneni, K. (2019). Modeling tabular data using conditional GAN. *Advances in Neural Information Processing Systems, 32*.

Yoon, J., Jarrett, D., & van der Schaar, M. (2019). Time-series generative adversarial networks. *Advances in Neural Information Processing Systems, 32*.

Zhao, Z., Kunar, A., Birke, R., Van der Scheer, H., & Chen, L. Y. (2024). CTAB-GAN+: Enhancing tabular data synthesis. *Frontiers in Big Data, 6*, 1296508.

Aggarwal, A., Mittal, M., & Battineni, G. (2021). Generative adversarial networks: An overview of theory and applications. *International Journal of Information Management Data Insights, 1*(1), 100004.

Ajay, A., Du, Y., Gupta, A., Tenenbaum, J., Jaakkola, T., & Agrawal, P. (2022). Is conditional generative modeling all you need for decision-making? *arXiv*. https://arxiv.org/abs/2211.15657

Andrew, G., Thakkar, O., McMahan, B., & Ramaswamy, S. (2021). Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems, 34*, 17455–17466.

Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. MIT Press.

Bedi, P., Goyal, S. B., & Kumar, J. (2020). Basic structure on artificial intelligence: A revolution in risk management and compliance. In *Proceedings of the 3rd International Conference on Intelligent Sustainable Systems* (pp. 570–576). IEEE.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., … Liang, P. (2021). On the opportunities and risks of foundation models. *arXiv*. https://arxiv.org/abs/2108.07258

Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., & Zhang, C. (2022). Quantifying memorization across neural language models. In *International Conference on Learning Representations (ICLR)*.

Chen, J., Liu, J., Liang, Y., & Zhou, M. (2026a). HeteroGCL: A heterogeneous graph contrastive learning framework for scalable and sustainable cryptocurrency AML. *Applied Sciences, 16*(6), 2860.

Chen, J., Liu, J., Liang, Y., & Zhou, M. (2026b). KE-MLLM: A knowledge-enhanced multi-sensor learning framework for explainable fake review detection. *Applied Sciences, 16*(6), 2909.

Dwork, C. (2025). Differential privacy. In *Encyclopedia of Cryptography, Security and Privacy* (pp. 649–652). Springer.

Figueira, A., & Vaz, B. (2022). Survey on synthetic data generation, evaluation methods and GANs. *Mathematics, 10*(15), 2733.

Giomi, M., Boenisch, F., Wehmeyer, C., & Tasnádi, B. (2022). A unified framework for quantifying privacy risk in synthetic data. *arXiv*. https://arxiv.org/abs/2211.10459

Grinsztajn, L., Oyallon, E., & Varoquaux, G. (2022). Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems, 35*, 507–520.

Gui, J., Sun, Z., Wen, Y., Tao, D., & Ye, J. (2021). A review on generative adversarial networks: Algorithms, theory, and applications. *IEEE Transactions on Knowledge and Data Engineering, 35*(4), 3313–3332.

He, Z., Zhang, T., & Lee, R. B. (2019). Model inversion attacks against collaborative inference. In *Proceedings of the Annual Computer Security Applications Conference* (pp. 148–162).

Ho, J., & Salimans, T. (2022). Classifier-free diffusion guidance. *arXiv*. https://arxiv.org/abs/2207.12598

Hu, H., Salcic, Z., Sun, L., Dobbie, G., Yu, P. S., & Zhang, X. (2022). Membership inference attacks on machine learning: A survey. *ACM Computing Surveys, 54*(11s), 1–37.

Jayaraman, B., & Evans, D. (2019). Evaluating differentially private machine learning in practice. In *USENIX Security Symposium* (pp. 1895–1912).

Lebichot, B., Le Borgne, Y. A., He-Guelton, L., Oblé, F., & Bontempi, G. (2019). Deep learning domain adaptation techniques for credit card fraud detection. In *INNS Big Data and Deep Learning Conference* (pp. 78–88). Springer.

Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. *Philosophical Transactions of the Royal Society A, 379*(2194).

Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. *arXiv*. https://arxiv.org/abs/2304.07619

Padhi, I., Schiff, Y., Melnyk, I., Rigotti, M., Mroueh, Y., Dognin, P., … Altman, E. (2021). Tabular transformers for modeling multivariate time series. In *IEEE ICASSP* (pp. 3565–3569). IEEE.

Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research, 22*(57), 1–64.

Platzer, M., & Reutterer, T. (2021). Holdout-based empirical assessment of mixed-type synthetic data. *Frontiers in Big Data, 4*, 679939.

Rajabi, A., & Garibay, O. O. (2022). TabFairGAN: Fair tabular data generation with generative adversarial networks. *Machine Learning and Knowledge Extraction, 4*(2), 488–501.

Salem, A., Bhattacharya, A., Backes, M., Fritz, M., & Zhang, Y. (2020). Updates-Leak: Data set inference and reconstruction attacks in online learning. In *USENIX Security Symposium* (pp. 1291–1308).

Schmidt, M., & Simic, M. (2019). Normalizing flows for novelty detection in industrial time series data. *arXiv*. https://arxiv.org/abs/1906.06904

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE, 109*(5), 612–634.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2020). Score-based generative modeling through stochastic differential equations. *arXiv*. https://arxiv.org/abs/2011.13456

Stadler, T., Oprisanu, B., & Troncoso, C. (2022). Synthetic data—Anonymisation groundhog day. In *USENIX Security Symposium* (pp. 1451–1468).