Proceedings of the 5th International Conference on Financial Technology and Business Analysis (ICFTBA 2026)
DOI: 10.54254/2754-1169/2026.39288

A Review of Machine Learning-Enhanced Monte Carlo Methods for Financial Risk Management

Fanxiang Mo

Central University of Finance and Economics, Beijing, China
2025310478@email.cufe.edu.cn

Abstract. Monte Carlo simulation has been a standard tool in financial risk management for decades, yet it converges slowly and becomes expensive when models are high-dimensional, when risk must be measured in real time, or when extreme market states matter. Machine learning has entered this workflow in three main ways. Deep importance sampling trains a neural network to learn the measure change that makes rare tail events more likely under the sampler, so the variance of value-at-risk estimates drops sharply for a fixed number of paths. Generative models, Quant GANs among them, synthesize extreme yet plausible market states; because the synthetic paths reproduce heavy tails and volatility clustering, they ease the chronic shortage of tail data. A third line replaces expensive pricing or loss computations with a trained neural surrogate that answers within milliseconds, a good fit for tasks that need repeated valuation. This study reviews the methods behind these three lines and their applications, including value-at-risk and expected shortfall estimation, stress scenario generation, credit portfolio loss simulation, and deep hedging. It also weighs the open problems. Black-box models remain difficult to explain to regulators under FRTB, tail regions are data-poor, and no accepted validation standard has emerged. This study concludes by arguing for Bayesian uncertainty quantification inside the simulation loop, validation procedures that regulators can accept, and evaluation frameworks that price the new model risk.

Keywords: Monte Carlo simulation, machine learning, variance reduction, generative models, surrogate models

1. Introduction

Financial institutions have relied on Monte Carlo simulation since the mid-twentieth century. The method estimates value at risk (VaR) and expected shortfall (ES), simulates credit portfolios, prices derivatives, and supports stress testing. Its convergence error is O(N^{-1/2}) and, conveniently, does not grow with dimension. The catch is speed: one extra digit of precision demands roughly one hundred times as many paths, so real-time computation becomes costly [1]. Meanwhile the products themselves grow more complex, and the Basel Fundamental Review of the Trading Book (FRTB) has replaced VaR with ES while tightening stress-testing requirements [2]. Traditional Monte Carlo therefore runs into high computational cost, slow convergence, the curse of dimensionality, and static distributional assumptions that fit modern markets poorly. This study asks how machine

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2026.39288

learning can ease these limits and let Monte Carlo methods serve risk management and regulation in high-dimensional, real-time, and extreme-market settings.

Three strands of research speak to this question. Neural networks can learn changes of measure that cut estimation variance; unlike conventional techniques that lean on prior knowledge of the problem structure [3], deep importance sampling reduces variance severalfold. Generative models can augment historical data for scenario construction, and the Quant GAN of Wiese et al. [4] reproduces the heavy tails and volatility clustering of financial time series. Deep networks can also approximate expensive pricing and hedging functions, as in the deep-hedging framework of Buehler et al. [5] and the equal-risk pricing of Carbonneau and Godin [6], both of which soften the curse of dimensionality. What the literature still lacks is synthesis. The work remains technology-driven and scattered across applications, and the model risks that machine learning introduces have drawn little systematic analysis. Because these methods exchange potential model bias for computational efficiency, this study takes up three questions: (1) What mechanisms drive the three approaches, and how do they complement one another? (2) How well do they perform in market risk measurement, credit portfolio simulation, and derivative pricing and hedging? (3) How should the resulting model risks be assessed and governed, and how should institutions weigh them against efficiency gains?

The contributions are of two kinds. Conceptually, this study gathers three previously separate technical approaches into one analytical framework, explains their mechanisms and complementary roles, and proposes a two-dimensional evaluation that accounts for efficiency gains and newly introduced risks alike. Practically, it gives financial institutions a basis for choosing among these technologies and for testing them against FRTB model-validation requirements.

2. Fusion paradigms of machine learning and Monte Carlo simulation

Machine learning enters Monte Carlo simulation in three ways: variance reduction, generative scenario construction, and surrogate modeling. Each drops a learnable component into a different stage of the usual pipeline of random sampling, path valuation, and statistical aggregation. Variance-reduction methods reshape how paths are drawn so that estimates carry less noise. Generative methods widen the scenario set, above all in extreme market states. Surrogates stand in for expensive calculations when valuation must repeat. Since the three act on different stages, nothing keeps a single pipeline from using all of them together. The remainder of this section surveys the representative methods and applications of each approach.

2.1. Variance reduction paradigm

Deep importance sampling is the clearest case of machine learning-based variance reduction. The network learns a change of measure under which rare tail events appear more often; for a fixed budget of paths, estimation variance then falls. The idea extends the classical Monte Carlo pricing framework of Boyle [3]. Where conventional methods need analytical approximations, a network can parameterize high-dimensional, nonlinear changes of measure directly, which suits the dependence structures of multi-asset portfolios far better.

Training concentrates sampling weight on the regions that contribute most to the risk measure, and that concentration is what minimizes the estimator's variance. In reported experiments the VaR estimation error falls to about a quarter of what standard Monte Carlo delivers, and the number of paths needed to hit a given regulatory accuracy target drops by more than an order of magnitude. Market-risk estimation of VaR and ES is the primary use, though the method also fits low-frequency, high-loss scenarios in credit portfolios.

DOI: 10.54254/2754-1169/2026.39288

2.2. Generative scenario construction paradigm

Generative models supply what history cannot: extreme yet plausible market states. Tail observations are scarce by nature, and any sample window may simply miss the crises an institution fears. The Quant GAN of Wiese et al. [4] reproduces the heavy tails, volatility clustering, and leverage effects of multi-asset returns, which gives market-risk tail scenarios an empirical footing. Huang et al. [7] went further with a regime-specific conditional GAN that generates financial time series under different market regimes. Rizzato et al. [8] batch-generated synthetic financial scenarios with GANs and verified that the synthetic data matched the key statistical features of the observed data.

These models also dispense with prespecified parametric distributions. They learn cross-asset dependence from historical data instead, and the joint scenarios that come out often resemble observed market behavior more closely than a fitted parametric model would. Vuletić and Cont [9] put conditional generative models to work in data-driven hedging; their volatility-surface-based network, trained on SPX implied-volatility data, generated arbitrage-free surface scenarios that supported hedging strategies outperforming conventional Delta and Delta-Vega benchmarks over a multi-year out-of-sample period without retraining. Evidence of that kind suggests generative scenarios can hold up over the extended horizons that stress testing and regulatory scenario design demand. Conditional models set market-factor states explicitly, so the extreme scenarios they produce are reproducible and traceable while remaining internally consistent. Where a Monte Carlo application needs many paths, such models generate high-dimensional scenarios at scale and complement conventional stochastic simulation.

2.3. Surrogate model paradigm

The third approach approximates expensive probabilistic calculations and pricing functions with neural networks or deep kernel methods. After training, a surrogate answers in milliseconds what previously took far longer. The recipe is straightforward: train a lightweight network on a limited set of high-fidelity simulation results, then let it stand in for the original computation. Training cost amortizes across the many evaluations that follow, which is exactly why surrogates shine in repeated valuation and real-time settings.

Credit portfolio models pose high-dimensional integration over correlated defaults. Neural-network surrogates take over the expensive calculation of conditional default probabilities and exposures, and large-scale credit valuation adjustment (CVA) estimation and dynamic portfolio-level risk monitoring then fit within operational time limits. The practical effect is more timely credit-risk management.

The pricing side has its own evidence. Horvath et al. [10] trained a network that calibrates the entire implied-volatility surface of a rough-volatility model in milliseconds, roughly three orders of magnitude faster than conventional numerical optimization. Zhuang et al. [11] paired Gaussian processes with deep kernel learning to price high-dimensional American options, with uncertainty estimates attached. Deep hedging takes a different route altogether: Buehler et al. [5] set risk-neutral pricing aside and minimized hedging risk directly, under transaction costs and whatever other constraints apply. As those costs grow, the volatility of hedged profit and loss in their results stays stable and stays below the conventional Delta and Delta-Vega benchmarks.

None of this is entirely new. Least-squares Monte Carlo, in which Longstaff and Schwartz [12] estimated continuation values by regression, is an early precursor of the same idea. Anderson and

DOI: 10.54254/2754-1169/2026.39288

Ulrych [13] and Gambara et al. [14] later showed that deep networks and other machine-learning regressors speed up the pricing of American and path-dependent contracts at little cost in accuracy.

3. Challenges and future directions

Efficiency, however, has a price. Machine learning introduces model risks of its own, and answering them calls for methodological and regulatory work alike.

3.1. Interpretability and regulatory compliance

Interpretability comes first among the difficulties. FRTB and similar frameworks expect pricing and risk models to carry traceable error bounds and verifiable assumptions, and an opaque deep-learning model offers neither. What is needed are interpretability methods that expose the factors driving a network's decisions, whether in changes of measure, scenario generation, or pricing-function approximation, and that quantify predictive uncertainty at the same time. Model-risk management and regulatory review both depend on those capabilities. Shinozaki [15] put the situation plainly: deep hedging and deep calibration rest on a fairly mature methodology, yet explaining model behavior to regulators is still an open problem.

3.2. Data scarcity and validation standards

The second difficulty is data. Tail losses and extreme market events are rare by definition, so training data tends to leave parts of the relevant state space uncovered, and extrapolation turns unstable. Nor has any widely accepted validation standard emerged for generated scenarios; how to test them for distributional drift or arbitrage consistency remains unsettled.

Three priorities stand out for future work. Bayesian deep learning and related uncertainty-quantification techniques belong inside the simulation workflow itself, so that every risk estimate comes with a measure of its own uncertainty. Validation and transparency methods must satisfy regulators without giving up the computational gains that motivated the exercise in the first place. Institutions, for their part, need an evaluation framework that weighs the benefits of hybrid methods against their costs before any technology-selection decision. One further need cuts across all three: common benchmark datasets and error metrics. Absent those, results from different studies resist comparison, and standardization stalls.

4. Conclusion

Three conclusions emerge from this study. The first concerns fit: the three approaches complement rather than replace one another. Deep importance sampling adjusts the sampling measure and cuts the variance of VaR estimates. Generative models turn out reproducible, traceable, high-dimensional paths for stress testing and regulatory scenarios. Neural-network surrogates shoulder the computational burden of high-dimensional problems with rapid approximation. Since each acts on a different stage of the simulation workflow, combining them improves both the speed and the precision of risk measurement.

The second conclusion concerns what these methods really are: a trade of potential model bias for computational efficiency. How well the trade works depends on the coverage of the training data, on whether the model structure suits the problem, and on the quality of the uncertainty quantification. Judged on speed alone, such methods will always look good. A fair evaluation needs both dimensions, efficiency gained and risk introduced.

DOI: 10.54254/2754-1169/2026.39288

The third conclusion concerns the road ahead. Progress waits on three tasks: embedding Bayesian deep learning and uncertainty quantification in simulation workflows, working out validation and transparency methods that satisfy regulatory requirements, and settling on an evaluation framework that weighs benefit against cost. FRTB implementation is advancing and the underlying methods keep maturing; on both counts, wider use of machine learning-enhanced Monte Carlo in financial risk management looks likely.

References

[1] Cont, R. (2004). Monte Carlo methods in financial engineering. *Quantitative Finance, 4*(4), C46–C47.
[2] Grajales, C. A., & Medina Hurtado, S. (2023). Sensitivities-based method and expected shortfall for market risk under FRTB: Its impact on options risk capital. *Journal of Economics, Finance and Administrative Science, 28*(55), 96–115.
[3] Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics, 4*(3), 323–338.
[4] Wiese, M., Knobloch, R., Korn, R., & Kretschmer, P. (2020). Quant GANs: Deep generation of financial time series. *Quantitative Finance, 20*(9), 1419–1440.
[5] Buehler, H., Gonon, L., Teichmann, J., & Wood, B. (2019). Deep hedging. *Quantitative Finance, 19*(8), 1271–1291.
[6] Carbonneau, A., & Godin, F. (2021). Equal risk pricing of derivatives with deep hedging. *Quantitative Finance, 21*(4), 593–608.
[7] Huang, A., Khushi, M., & Suleiman, B. (2023). Regime-specific quant generative adversarial network: A conditional generative adversarial network for regime-specific deepfakes of financial time series. *Applied Sciences, 13*(19), 10639.
[8] Rizzato, M., Wallart, J., Geissler, C., Morizet, N., & Boumlaik, N. (2023). Generative adversarial networks applied to synthetic financial scenarios generation. *Physica A: Statistical Mechanics and its Applications, 623*, 128899.
[9] Vuletić, M., & Cont, R. (2024). VolGAN: A generative model for arbitrage-free implied volatility surfaces. *Applied Mathematical Finance, 31*(4), 203–238.
[10] Horvath, B., Muguruza, A., & Tomas, M. (2021). Deep learning volatility: A deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance, 21*(1), 11–27.
[11] Zhuang, J., Ding, D., Lu, W., Wu, X., & Yuan, G. (2025). A Gaussian process based method with deep kernel learning for pricing high-dimensional American options. *Computational Economics, 66*(5), 3687–3708.
[12] Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies, 14*(1), 113–147.
[13] Anderson, D., & Ulrych, U. (2023). Accelerated American option pricing with deep neural networks. *Quantitative Finance and Economics, 7*(2), 207–228.
[14] Gambara, M., Livieri, G., & Pallavicini, A. (2025). Machine-learning regression methods for American-style path-dependent contracts. *Quantitative Finance, 25*(6), 895–918.
[15] Shinozaki, Y. (2025). Deep learning in finance: A review of deep hedging and deep calibration techniques. *International Journal of Theoretical and Applied Finance, 28*(3–4), 1–44.