Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

The Evolution of Monte Carlo Simulation: From Traditional Methods to AI-Driven Financial Risk Modeling

NARESH SRITHAREN
Independent Researcher, USA
Corresponding Author: NARESH SRITHAREN, **E-mail**: sritharenn@gmail.com

| ABSTRACT
Monte Carlo Simulation has served as a fundamental methodology in financial risk management for decades, particularly in applications such as Interest Rate Risk in the Banking Book, market risk assessment, and regulatory stress testing frameworks. While traditional Monte Carlo approaches have proven valuable for probabilistic risk modeling, they face significant limitations, including reliance on static distributions, fixed correlation matrices, a lack of macroeconomic coherence in scenario generation, and substantial computational demands. The emergence of artificial intelligence and generative AI technologies presents a transformative opportunity to address these shortcomings while preserving the core probabilistic framework that makes Monte Carlo simulation effective. This article examines the comprehensive evolution of Monte Carlo simulation through AI integration, exploring how technologies such as Variational Autoencoders, Graph Neural Networks, Generative Adversarial Networks, and diffusion models enhance each phase of the simulation lifecycle from data preparation through governance. The transformation encompasses richer data ingestion through vector databases and natural language processing, neural density estimation for capturing complex distributions, dynamic correlation modeling that adapts to market regimes, generation of macroeconomically consistent stress scenarios, accelerated valuation through neural surrogate models, and enhanced explainability via SHAP values and large language model-generated reports. Through detailed examination of Interest Rate Risk in the Banking Book applications, particularly Economic Value of Equity analysis, this article demonstrates how AI-enhanced simulation produces results that are simultaneously more accurate, comprehensive, explainable, and aligned with regulatory expectations for model risk management and stress testing transparency, representing not a replacement but an intelligent augmentation of proven quantitative methods.

| KEYWORDS
Monte Carlo Simulation, Generative AI, Interest Rate Risk in the Banking Book, Financial Risk Management, Explainable Artificial Intelligence.

| ARTICLE INFORMATION
ACCEPTED: 01 November 2025 **PUBLISHED:** 21 November 2025 **DOI:** 10.32996/jcsts.2025.7.12.11

1. Introduction
Monte Carlo Simulation (MCS) has long served as a cornerstone methodology in financial risk management, particularly within Interest Rate Risk in the Banking Book (IRRBB), market risk assessment, and Comprehensive Capital Analysis and Review (CCAR) frameworks. Named after the famed casino in Monaco to reflect its reliance on randomness and probability, this technique estimates uncertain outcomes by generating thousands of possible scenarios through random sampling. In its traditional form, MCS employs probabilistic modeling to evaluate risk metrics such as Economic Value of Equity (EVE), Net Interest Income (NII), and Value at Risk (VaR) by defining problem parameters, identifying uncertain inputs, assigning probability distributions, and aggregating results across numerous iterations. Research examining IRRBB levels for Indian banks has demonstrated that interest rate risk exposure varies significantly across institutions, with traditional simulation approaches revealing that changes in interest rates directly impact both the economic value of banks' equity and their net interest income streams, as documented by Sathye and Sathye in their comprehensive analysis of Indian banking institutions [1].

The widespread adoption of Monte Carlo methods in banking has been driven by regulatory mandates and the need to quantify complex risk exposures under uncertainty. Financial institutions employ these simulations to model how interest rate movements affect their entire balance sheet, considering the repricing characteristics of assets and liabilities with different maturities and embedded options. The IRRBB framework requires banks to assess potential losses from parallel and non-parallel shifts in yield curves, with traditional Monte Carlo approaches generating multiple interest rate scenarios to estimate the distribution of possible outcomes. Studies of Indian banks have shown that IRRBB exposure levels depend heavily on the composition of the balance sheet, the maturity profile of assets and liabilities, and the behavioral assumptions embedded in deposit and prepayment models, with some institutions demonstrating higher vulnerability to rate shocks than others based on their asset-liability mismatches [1].

However, the financial landscape's increasing complexity—characterized by regime shifts, tail risk dependencies, and nonlinear relationships—has exposed critical limitations in conventional MCS approaches. The challenge of modeling correlations and dependencies between risk factors has become particularly acute, as traditional linear correlation measures fail to capture the complex dependency structures observed in financial markets. Research on copula selection for risk management has emphasized that the choice of dependency structure significantly impacts risk measures, with empirical evidence showing that tail dependencies during stress periods differ markedly from normal market conditions, as explored by Embrechts, McNeil, and Straumann in their foundational work on risk management methodologies [2]. Traditional methods that rely on static distributions and fixed correlation matrices struggle to capture these dynamic relationships and fail to reflect the macroeconomic realism necessary for comprehensive risk assessment. The computational resources required for traditional approaches further constrain the breadth and depth of scenario analysis that institutions can feasibly conduct.

The advent of artificial intelligence (AI) and generative AI (GenAI) technologies presents a transformative opportunity to address these shortcomings while preserving the fundamental probabilistic framework that makes Monte Carlo simulation valuable. This article examines how AI and GenAI enhance each phase of the Monte Carlo lifecycle, from data preparation through governance, ultimately creating a more robust, explainable, and dynamically responsive risk modeling ecosystem that addresses both the dependency modeling challenges identified in copula research and the balance sheet risk assessment imperatives demonstrated in banking sector studies.

2. Traditional Monte Carlo Simulation: Methodology and Limitations

The classical Monte Carlo simulation framework follows a systematic seven-phase lifecycle that has been refined over decades of application in financial risk management. Beginning with data preparation, analysts collect historical market data, including interest rates, spreads, volatilities, and correlations, using standard tools such as SQL, Python, and R. This foundational phase requires careful curation of time series data, often spanning multiple years to capture different market regimes and stress periods. The assumption-setting phase involves defining statistical distributions—typically Normal, Lognormal, or Student-t distributions—for each risk factor based on historical calibration. Monte Carlo simulation in financial engineering has evolved into a sophisticated computational technique that addresses complex valuation and risk management problems, particularly for derivative securities and portfolios with nonlinear payoffs where analytical solutions prove intractable, as extensively documented by Glasserman in his comprehensive treatment of Monte Carlo methods applied to financial problems [3]. Correlation modeling then constructs covariance matrices to capture interdependencies among risk factors, commonly employing Pearson correlation coefficients and Cholesky decomposition for multivariate simulation. The Cholesky decomposition method transforms independent random variables into correlated ones by factorizing the positive-definite covariance matrix, enabling the generation of correlated scenarios that preserve the specified dependency structure across multiple risk factors simultaneously.

The core of the simulation occurs during scenario generation, where pseudorandom or quasi-random number generators produce thousands of potential paths for each risk factor. Quasi-random sequences like Sobol or Halton sequences offer improved convergence properties compared to pseudorandom generators, particularly for high-dimensional problems common in portfolio risk assessment, where dozens or even hundreds of risk factors must be simulated simultaneously. Each generated scenario feeds into valuation models that reprice portfolios using discounted cash flow methodologies, calculating metrics like EVE and NII under various conditions. The aggregation phase synthesizes these results into summary statistics, computing quantile-based risk measures such as VaR and Expected Shortfall. The application of Monte Carlo methods to financial engineering problems requires careful attention to variance reduction techniques, path generation algorithms, and computational efficiency considerations, particularly when dealing with American-style options or portfolios containing multiple exotic derivatives that demand intensive calculation for each simulated path [3]. Finally, validation and governance processes ensure model accuracy through backtesting against historical outcomes and sensitivity analysis, adhering to regulatory frameworks like SR 11-7 for model risk management.

Despite its widespread adoption, traditional MCS faces significant limitations that have become increasingly apparent as financial markets have grown more complex and interconnected. The assumption of static distributions fails to capture regime changes and structural breaks in financial markets, where volatility patterns and return characteristics can shift dramatically during crisis periods. Fixed correlation matrices cannot adequately represent time-varying dependencies or tail co-movements during stress periods, a phenomenon that has gained renewed attention with the emergence of artificial intelligence tools in financial analysis. Recent research examining correlation pitfalls has highlighted how misunderstanding correlation concepts can lead to flawed risk assessments, particularly when practitioners fail to distinguish between correlation and causation or when they apply correlation measures inappropriately to non-stationary data, as demonstrated in studies analyzing common statistical misconceptions that persist even among sophisticated users of quantitative methods [4]. Scenario generation often lacks macroeconomic coherence, producing paths that may be mathematically valid but economically implausible. Furthermore, the computational intensity of running thousands or millions of scenarios creates practical constraints, while the static nature of calibrated models means they quickly become outdated as market conditions evolve. These limitations have motivated the search for more adaptive, intelligent approaches to probabilistic risk modeling.

Phase	Traditional Approach	AI/GenAI Enhancement	Key Improvement Metric
Data Preparation	Historical time series only (rates, spreads, volatilities)	Structured + unstructured data (news, filings, text) via AI data fabrics	Data richness: 2-3x more information sources
Assumption Setting	Parametric distributions (Normal, Lognormal, Student-t)	Neural density estimators, VAEs, learn true distributions	Distribution accuracy: Captures multimodal, fat-tailed behavior
Correlation Modeling	Static covariance matrix (Pearson correlation, Cholesky)	Graph Neural Networks model dynamic dependencies	Correlation adaptability: Regime-dependent relationships
Scenario Generation	Pseudorandom/quasi-random draws (Sobol, Halton sequences)	GANs and diffusion models generate macro-consistent scenarios	Economic realism: Plausible, stress-aligned paths
Valuation/Projection	Formula-based DCF pricing models	Neural surrogates, RL simulators approximate pricing	Computation speed: 10-100x faster
Aggregation	Quantile metrics (VaR, Expected Shortfall)	Explainable AI (SHAP) for driver attribution	Insight depth: Traces risk to specific drivers
Validation/Governance	Static backtesting, periodic recalibration	Continuous monitoring, drift detection, and audit trails	Governance quality: Real-time model validation

Table 1: Comparison of Traditional vs. AI-Enhanced Monte Carlo Simulation Phases [3, 4]

3. AI and GenAI Transformation Across the Simulation Lifecycle

The integration of AI and GenAI technologies fundamentally transforms each phase of the Monte Carlo simulation process, addressing traditional limitations while introducing new capabilities that reshape how financial institutions approach risk modeling and scenario analysis. In data preparation, AI-powered data fabrics and vector databases enable the integration of both structured time series and unstructured data sources—including news feeds, regulatory filings, and textual information—creating a richer, real-time data environment. This expansion beyond purely numerical historical data allows models to incorporate forward-looking information and sentiment signals that may presage market movements. Natural language processing techniques can extract risk signals from Federal Reserve communications, earnings call transcripts, and regulatory announcements, transforming qualitative information into quantitative inputs for risk models. Machine learning algorithms can process these heterogeneous data streams continuously, identifying emerging patterns and regime shifts that traditional models based solely on historical price data would miss.

Assumption setting evolves from parametric distribution fitting to neural density estimation, where Variational Autoencoders (VAEs) and other deep learning architectures learn the true underlying distributions of risk factors directly from data. These methods can capture nonlinear relationships, multimodal distributions, and fat-tailed behavior that parametric assumptions often miss. Traditional approaches that assume returns follow normal or log-normal distributions systematically underestimate the probability of extreme events, whereas neural density estimators can learn arbitrary probability distributions that better reflect the empirical characteristics of financial data. Correlation modeling advances through Graph Neural Networks (GNNs), which model dynamic dependencies and regime-dependent correlations rather than assuming static relationships. This allows the simulation to adapt correlation structures based on market conditions, capturing the tendency for correlations to increase during stress periods. Research on deep learning for portfolio optimization has demonstrated that neural network architectures can effectively learn complex patterns in financial data and optimize asset allocation strategies, with studies showing that deep learning models incorporating Long Short-Term Memory networks and Convolutional Neural Networks can process historical price data to generate portfolio weights that adapt to changing market conditions and outperform traditional mean-variance optimization approaches in various market scenarios [5].

Perhaps most significantly, scenario generation transforms through the application of Generative Adversarial Networks (GANs) and diffusion models. Rather than simple random draws from predetermined distributions, these GenAI models produce realistic, macroeconomically consistent stress scenarios that can be conditioned on specific policy parameters or economic assumptions. GANs consist of two neural networks—a generator that creates synthetic scenarios and a discriminator that distinguishes real from generated data—trained in an adversarial fashion until the generator produces scenarios indistinguishable from historical observations. These generative models can be conditioned on macroeconomic variables such as GDP growth, unemployment rates, or central bank policy stances, ensuring that generated interest rate paths remain consistent with broader economic scenarios. In the valuation phase, neural surrogate models and reinforcement learning simulators approximate complex pricing functions with dramatic computational speedups, enabling more sophisticated modeling of nonlinear instruments. Neural networks have proven particularly effective at solving partial differential equations that govern derivative pricing and risk evolution, with deep learning algorithms demonstrating the capability to approximate solutions to complex PDEs that arise in financial mathematics, including the Black-Scholes equation and its generalizations, thereby enabling rapid valuation of options and structured products across thousands of scenarios without requiring computationally expensive numerical methods for each iteration [6].

Aggregation and insight generation benefit from Explainable AI techniques such as SHAP values, which trace portfolio sensitivity back to specific risk drivers, while large language models generate natural language summaries of results. Finally, governance evolves from static backtesting to continuous model monitoring with drift detection, maintaining audit trails through model registries and explainability reports that satisfy regulatory requirements.

Simulation Phase	Traditional Technology	AI/GenAI Technology	Specific AI Method	Primary Benefit
Data Preparation	SQL, Python, R for structured data	AI data fabrics, Vector databases	Natural Language Processing	Real-time multi-source integration
Assumption Setting	Parametric fitting (Normal, Lognormal)	Neural density estimation	Variational Autoencoders (VAEs)	Captures nonlinear, multimodal distributions
Correlation Modeling	Pearson correlation, Static matrices	Dynamic dependency networks	Graph Neural Networks (GNNs)	Regime-dependent correlations
Scenario Generation	Pseudorandom generators (Sobol, Halton)	Generative models	GANs, Diffusion models	Macro-consistent stress scenarios
Valuation/Pricing	Discounted cash flow formulas	Neural approximators	Neural surrogates, PDE solvers	10-100x computational speedup
Behavioral Modeling	Fixed prepayment curves	Adaptive decision models	Reinforcement Learning	Context-dependent responses
Aggregation/Insights	VaR, Expected Shortfall calculations	Attribution analysis	SHAP values, XAI techniques	Driver-level sensitivity tracing
Reporting	Manual report generation	Automated narrative generation	Large Language Models (LLMs)	Natural language summaries
Governance	Periodic backtesting	Continuous monitoring	Drift detection algorithms	Real-time model validation

Table 2: AI Technologies Applied Across Monte Carlo Simulation Phases [5, 6]

4. Architectural Framework for GenAI-Enabled Simulation

The implementation of GenAI-enhanced Monte Carlo simulation requires a comprehensive architectural framework that integrates multiple AI components into a cohesive pipeline designed to address the complex requirements of modern financial risk management. The ingestion layer continuously streams market data, macroeconomic indicators, and textual signals from diverse sources—including interest rates, credit spreads, Federal Reserve releases, and corporate earnings reports—into vector databases optimized for both structured and unstructured data. This real-time data infrastructure ensures that models always operate on current information rather than becoming stale between recalibration cycles. Modern data architectures leverage distributed computing frameworks and cloud-based storage solutions to handle the massive volumes of heterogeneous data generated by financial markets, with implementations processing millions of price updates, thousands of news articles, and hundreds of economic releases daily. Vector databases employ embedding techniques that transform both numerical time series and textual documents into high-dimensional vector representations, enabling efficient similarity search and retrieval of relevant historical patterns that inform current risk assessments.

The learning layer trains multiple AI models in parallel, with VAEs or GANs learning the joint distribution of risk factors from historical and current data, while GNNs capture the dynamic correlation network and its evolution across different market regimes. These models update continuously or semi-continuously as new data arrives, maintaining relevance to current market conditions. The training process employs techniques such as transfer learning and incremental learning to efficiently incorporate new observations without requiring complete retraining from scratch. Research on Machine Learning Operations has established comprehensive frameworks for managing the end-to-end lifecycle of ML systems in production environments, emphasizing that

MLOps encompasses workflow orchestration, model training automation, continuous integration and deployment pipelines, model versioning, monitoring systems, and governance structures that ensure models remain performant and compliant throughout their operational lifecycle, with studies demonstrating that organizations implementing mature MLOps practices achieve faster model deployment cycles and improved model reliability compared to those relying on ad-hoc processes [7]. The generation layer employs conditional diffusion models or GANs to create plausible macroeconomic stress scenarios, with the conditioning mechanism allowing risk managers to align scenarios with specific regulatory stress parameters or hypothetical policy changes from central banks.

The simulation layer feeds these AI-generated scenarios into pricing models, which may themselves be AI-based surrogate models trained to approximate complex valuation functions. Reinforcement learning agents can model dynamic customer behavior—such as prepayment rates and deposit flows—that respond to changing conditions rather than following fixed assumptions. The explanation layer applies XAI techniques to decompose results, identifying which input shocks drive portfolio sensitivity and risk metrics. SHAP values and related attribution methods provide quantitative measures of feature importance, enabling risk managers to understand not just what the model predicts but why it produces particular outputs. Research on Explainable Artificial Intelligence has identified fundamental concepts and challenges in making AI systems interpretable, noting that XAI encompasses multiple approaches including feature importance methods, rule extraction techniques, attention mechanisms, and counterfactual explanations, with particular emphasis on the healthcare domain demonstrating that explainability proves crucial for building trust with domain experts, facilitating regulatory compliance, and enabling practitioners to validate that model predictions align with established medical knowledge rather than spurious correlations, principles that extend directly to financial risk management where stakeholders similarly require transparent understanding of model reasoning to satisfy governance requirements [8]. LLM-based report generators synthesize these technical attributions into natural language narratives accessible to both technical and non-technical stakeholders. Finally, the governance layer maintains a model registry using platforms like MLflow or Vertex AI, tracks model performance and drift, generates explainability reports for auditors, and incorporates human validation loops to ensure AI-generated scenarios and valuations remain reasonable and defensible.

Architecture Layer	Primary Function	Key Technologies	Data Processing Capability	Output/Deliverable
Ingestion Layer	Real-time data streaming	Vector databases, Embedding techniques	Millions of price updates, thousands of news articles, hundreds of economic releases daily	Unified structured + unstructured data repository
Learning Layer	Model training and updating	VAEs, GANs, GNNs, Transfer learning	Continuous/semi-continuous learning from new data	Trained distribution models, dynamic correlation networks
Generation Layer	Scenario creation	Conditional diffusion models, GANs	Regulatory stress parameters, policy constraints	Macro-consistent stress scenarios
Simulation Layer	Portfolio valuation	Neural surrogates, Reinforcement Learning agents	Complex pricing functions, behavioral dynamics	EVE, NII, risk metrics across scenarios
Explanation Layer	Result decomposition	SHAP values, XAI techniques, and Attribution methods	Feature importance quantification	Driver-level sensitivity analysis
Reporting Layer	Narrative generation	Large Language Models (LLMs)	Technical-to-natural language translation	Stakeholder-accessible reports
Governance Layer	Model oversight	MLflow, Vertex AI, Drift detection algorithms	Performance tracking, audit trail maintenance	Compliance reports, validation documentation

Table 3: GenAI-Enhanced Monte Carlo Simulation Architecture - Layer-by-Layer Framework [7, 8]

5. Application Case Study: IRRBB and Economic Value of Equity Analysis

The practical benefits of GenAI-enhanced simulation become apparent when comparing traditional and AI-driven approaches to Interest Rate Risk in the Banking Book analysis, specifically in calculating EVE under interest rate stress. In the traditional approach, scenario inputs consist of randomly generated interest rate paths drawn from calibrated parametric distributions, with correlations captured by a fixed covariance matrix. Customer behavior follows predetermined prepayment curves based on historical analysis, and the EVE output reports mean values alongside extreme quantiles such as the percentile loss thresholds that regulators require banks to monitor. Results are typically presented in static stress tables with limited insight into the underlying drivers of risk. Research on understanding and managing interest rate risk at banks has emphasized that effective IRRBB management requires comprehensive frameworks that measure both earnings-based metrics like NII and economic value measures like EVE, with studies highlighting that banks must consider multiple dimensions of interest rate risk including repricing risk arising from timing differences in the maturity and repricing of assets and liabilities, basis risk from imperfect correlation between different rate indices, yield curve risk from non-parallel shifts in the term structure, and embedded option risk from prepayments and early withdrawals that customers exercise in response to rate movements [9].

The GenAI transformation enhances each component of this analysis through the application of advanced machine learning architectures and generative modeling techniques. Scenario inputs now consist of macro-consistent rate paths generated by conditional diffusion models that respect economic relationships and policy constraints, ensuring that simulated interest rate trajectories remain aligned with plausible macroeconomic conditions rather than producing mathematically valid but economically implausible combinations. Rate correlations emerge from GNN models that adapt to current regime characteristics rather than assuming historical relationships will persist indefinitely into the future. Customer behavior is modeled by AI systems trained on both historical data and synthetic scenarios, capturing how prepayments and deposit flows respond to rate changes in nonlinear, context-dependent ways that reflect the heterogeneity of customer populations. The EVE output expands from simple summary statistics to full probability surfaces with explainable driver attribution, identifying which specific rate movements, curve shapes, or behavioral responses contribute most significantly to potential losses. Studies examining interest rate risk management frameworks have documented that banks employing sophisticated modeling techniques that account for behavioral dynamics and non-parallel yield curve shifts achieve more accurate risk assessments compared to institutions relying solely on standard regulatory shock scenarios, enabling better-informed capital allocation and hedging decisions [9].

Perhaps most valuably, reporting evolves from static tables to auto-generated IRRBB stress reports with sensitivity narratives produced by LLMs. These reports explain in natural language why certain scenarios produce particular EVE impacts, which portfolio positions drive vulnerability, and how different risk factors interact to amplify or mitigate losses under stress. Research reviewing deep learning developments has highlighted the transformative potential of neural network architectures across diverse application domains, noting that deep learning methods have demonstrated remarkable success in pattern recognition, time series forecasting, and complex decision-making tasks, with the technology's continued evolution driven by advances in network architectures, optimization algorithms, and the availability of large-scale training datasets that enable models to learn increasingly sophisticated representations of underlying data structures [10]. This transformation does not abandon the Monte Carlo framework but rather elevates it, producing results that are simultaneously more accurate, more comprehensive, more explainable, and more aligned with regulatory expectations for model risk management and stress testing transparency.

Risk Dimension	Definition	Traditional Measurement Approach	GenAI-Enhanced Approach	Key Improvement
Repricing Risk	Timing differences in asset/liability maturity	Static gap analysis with fixed assumptions	Dynamic gap modeling with behavioral AI	Context-dependent repricing patterns
Basis Risk	Imperfect correlation between rate indices	Fixed correlation matrices	GNN-based dynamic correlation modeling	Regime-adaptive basis relationships
Yield Curve Risk	Non-parallel shifts in term structure	Standardized parallel shock scenarios	Conditional diffusion model scenarios	Macroeconomically consistent curve shifts
Embedded Option Risk	Customer prepayments and early withdrawals	Predetermined prepayment curves	RL-based behavioral models with synthetic data	Nonlinear, heterogeneous responses
EVE Measurement	Economic value impact of rate changes	Mean + percentile quantiles (e.g., 99th)	Full probability surface with driver attribution	Comprehensive distribution understanding
NII Measurement	Earnings impact over the time horizon	Simple projection with fixed behaviors	AI-driven multi-scenario projections	Forward-looking earnings volatility
Scenario Generation	Interest rate path creation	Random draws from parametric distributions	GANs/diffusion models with policy constraints	Economic plausibility and coherence
Risk Attribution	Understanding loss drivers	Limited sensitivity analysis	SHAP values for driver decomposition	Granular factor-level insights
Reporting Format	Communication to stakeholders	Static stress tables	LLM-generated natural language narratives	Stakeholder accessibility and clarity

Table 4: IRRBB Risk Dimensions - Traditional vs. GenAI-Enhanced Measurement [9, 10]

6. Conclusion

The evolution of Monte Carlo simulation through artificial intelligence and generative AI integration represents a fundamental advancement in financial risk management methodology that addresses critical limitations of traditional approaches while preserving the probabilistic rigor that has made Monte Carlo methods indispensable to the banking industry. This transformation extends across the entire simulation lifecycle, from data preparation enhanced by vector databases and natural language processing that incorporate diverse information sources, through assumption setting powered by neural density estimators that capture complex, multimodal distributions, to correlation modeling via Graph Neural Networks that adapt dynamically to changing market regimes. The application of Generative Adversarial Networks and diffusion models to scenario generation produces macroeconomically coherent stress paths that respect fundamental economic relationships and policy constraints, overcoming the implausible combinations often generated by traditional random sampling approaches. Neural surrogate models and reinforcement learning agents dramatically accelerate computation while enabling more sophisticated modeling of nonlinear instruments and context-dependent customer behaviors, expanding the feasible scope of risk analysis from thousands to hundreds of thousands of scenarios within operational timeframes. The integration of Explainable AI techniques, such as SHAP values, provides unprecedented transparency in risk attribution, tracing portfolio sensitivity back to specific drivers, and enabling large language models to generate natural language narratives that communicate complex technical findings to diverse stakeholders. The comprehensive architectural framework presented, spanning ingestion, learning, generation, simulation, explanation, reporting, and governance layers, demonstrates how mature Machine Learning Operations practices can ensure that AI-enhanced models remain performant, compliant, and defensible throughout their operational lifecycle. The Interest Rate Risk in the Banking Book case article demonstrates these advantages tangibly, demonstrating how GenAI transformation lifts Economic Value of Equity analysis from rigid stress tables to dynamic, interpretable probability surfaces that get at repricing risk, basis risk, yield curve risk, and embedded option risk more precisely and with a finer granularity than traditional approaches allow. This integration of human capability, proven quantitative methods, and advanced AI technologies is a prime example of how financial risk management innovation advances best through smart augmentation, not complete displacement, to build more solid, transparent, and responsive frameworks to the intricacies of current financial markets that continue to uphold the rigor of governance required by regulators and internal constituents alike for high-risk capital deployment and strategic decisions.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.

References

[1] Dominic K et al., (2023) Machine Learning Operations (MLOps): Overview, Definition, and Architecture, ResearchGate, January 2023. [Online]. Available: https://www.researchgate.net/publication/369566050_Machine_Learning_Operations_MLOps_Overview_Definition_and_Architecture

[2] Eric K et al., (2026) Selecting Copulas for Risk Management, *Journal of Banking & Finance,* ResearchGate. May 2026. [Online]. Available: https://www.researchgate.net/publication/222433063_Selecting_Copulas_for_Risk_Management

[3] Marius H, (2023) Correlation Pitfalls with ChatGPT: Would You Fall for Them? ResearchGate, June 2023. [Online]. Available: https://www.researchgate.net/publication/371800623_Correlation_Pitfalls_with_ChatGPT_Would_You_Fall_for_Them

[4] Meera S, (2012) Interest Rate Risk in the Banking Book (IRRBB): Levels for Indian Banks, ResearchGate, May 2012. [Online]. Available: https://www.researchgate.net/publication/256019537_Interest_Rate_Risk_in_the_Banking_Book_IRRBB_Levels_for_Indian_Banks

[5] Nan C & Jeff H V, (2008) Monte Carlo simulation in financial engineering, ResearchGate, January 2008. [Online]. Available: https://www.researchgate.net/publication/4306230_Monte_Carlo_simulation_in_financial_engineering

[6] Soumaya N et al., (2024) Utilizing deep learning algorithms for the resolution of partial differential equations, ResearchGate, December 2024. [Online]. Available: https://www.researchgate.net/publication/386099481_Utilizing_deep_learning_algorithms_for_the_resolution_of_partial_differential_equations

[7] Tim H, (2023) Explainable Artificial Intelligence (XAI): Concepts and Challenges in Healthcare, ResearchGate, August 2023. [Online]. Available: https://www.researchgate.net/publication/373079761_Explainable_Artificial_Intelligence_XAI_Concepts_and_Challenges_in_Healthcare

[8] Viral V. A, (2018) Understanding and managing interest rate risk at banks, Federal Reserve Bank of Cleveland and Federal Reserve Bank of New York, ResearchGate, July 2018. [Online]. Available: https://www.researchgate.net/publication/326233681_Understanding_and_managing_interest_rate_risk_at_banks

[9] Zhiying H et al., (2019) Deep learning review and discussion of its future development, ResearchGate, January 2019. [Online]. Available: https://www.researchgate.net/publication/332150242_Deep_learning_review_and_discussion_of_its_future_development

[10] Ziao Z et al., (2020) Deep Learning for Portfolio Optimization, ResearchGate, October 2020. [Online]. Available: https://www.researchgate.net/publication/346737516_Deep_Learning_for_Portfolio_Optimization