ICCAUA Proceedings Journal
Proceedings of the international conference of contemporary affairs in architecture and urbanism-ICCAUA
Volume 9 (December 2026), 2610121

Journal homepage: https://journal.iccaua.com/

DOI: https://doi.org/10.38027/ICCAUA2026EN0121

Rethinking Operational Risk in the AI Era: A Comparative Perspective

* [1] Chelh Sara , [2] Ababou Mariame , [3] El Gharbaoui Bouteina
[1 & 3]Management & Innovation Research Laboratory-MIRL, Department of Management, National School of Business and Management ENCG, University of Moulay Ismail, Meknes, Morocco
[2] Higher School of Technology, University Mohamed 1st Oujda, Nador, Morocco
[1] E-mail: s.chelh@umi.ac.ma , [2] E-mail: : m.ababou@ump.ac.ma , [3] E-mail: bouteina.elgharbaoui@gmail.com

Received: 26.04.2026
Revised: 24.06.2026
Accepted: 01.07.2026
Available online: 10.07.2026

This article has been selected and peer-reviewed for publication in this journal as part of the 9th International Conference of Contemporary Affairs in Architecture and Urbanism, held on 7–8 May 2026 in Istanbul, Türkiye.

Abstract

Operational risk has become increasingly important due to recurrent financial scandals and significant loss events, yet it remains comparatively underexplored in the literature. At the same time, artificial intelligence is gaining prominence as a transformative tool in risk management. This paper compares traditional parametric and machine-learning approaches for operational loss severity modeling using the SAS OpRisk database. The analysis combines descriptive examination of loss characteristics, parametric fitting with the Lognormal and Gamma distributions, and supervised prediction based on structured explanatory variables. The predictive framework includes Ridge regression, Decision Tree, Random Forest, Gradient Boosting, and a hybrid Ridge–Random Forest model. The findings suggest that operational loss severity reflects both linear and non-linear structures and that hybrid modeling can improve predictive performance. The research highlights the complementary roles of traditional parametric fitting and machine-learning methods in operational risk severity analysis
Keywords: Operational Risk; Artificial Intelligence (AI); Machine learning; Loss Prediction; Databases.

1. Introduction

Operational risk now occupies a central place in financial risk management due to its complexity, the diversity of its sources, and its capacity to generate substantial losses. Its multifaceted nature makes it both conceptually broad and methodologically difficult to quantify. Against this background, the introduction is organized into two subsections. The first discusses operational risk as a concept and a contemporary regulatory and managerial challenge, while the second turns to the specific issue of operational loss severity and the evolution of its modeling approaches, from traditional parametric methods to AI-based techniques.

1.1 Operational Risk: Concept, Scope, and Contemporary Relevance

Operational risk has evolved into a fundamental pillar of banking supervision and enterprise risk management. Under the Basel framework, operational risk encompasses losses stemming from inadequate or failed internal processes, people, systems, or external events (Basel Committee on Banking Supervision [BCBS], 2021; Chernobai et al., 2018). This comprehensive definition underscores the multidimensional and heterogeneous character of operational risk, which spans internal fraud, employment practices, client-related failures, business disruption, system failures, execution errors, and damage to physical assets (Cope et al., 2012; Dahen & Dionne, 2010; Moosa, 2007). Recent supervisory guidance has further integrated operational resilience as a complementary dimension, particularly concerning severe but plausible disruptions, including cyber-attacks, pandemics, and critical infrastructure failures (BCBS, 2021; Financial Stability Board, 2020).
Operational risk fundamentally differs from market and credit risk in that it does not originate from discrete financial exposures but is intrinsically embedded within institutional operations and organizational processes (Cruz et al., 2015). This embeddedness renders operational risk inherently more challenging to identify, classify, measure, and manage. The complexity is further compounded by the dual nature of operational losses: the coexistence of high-frequency, low-impact events alongside low-frequency, high-severity tail events, combined with intricate interactions among organizational culture, technological infrastructure, process design, and exogenous shocks (Moscadelli, 2004; Dutta & Perry, 2006; Cope et al., 2012). Consequently, operational risk management has ascended to a central position within supervisory frameworks and corporate governance structures. The Basel Committee's revised Principles for the Sound Management of Operational Risk and its Principles for Operational Resilience explicitly emphasize robust governance frameworks, resilient information and communication technology (ICT) infrastructures, comprehensive business continuity planning, and effective change management as foundational elements of sound operational risk practice (BCBS, 2021, 2024).

The salience of operational risk has been substantially amplified by accelerating digitalization and the pervasive adoption of advanced technologies within financial services. Empirical evidence demonstrates that digital transformation simultaneously expands institutional exposure to operational, cyber, and third-party risks while enabling sophisticated technological solutions for risk surveillance, early warning systems, and mitigation strategies (Aldasoro et al., 2020; Financial Stability Board, 2023). The recent proliferation of artificial intelligence (AI) and machine learning applications introduces a paradox: enhanced analytical capabilities and operational efficiencies coexist with emergent risks, including model opacity, algorithmic bias, data quality vulnerabilities, and novel operational failure modes. This dual-edged nature necessitates more sophisticated, dynamic, and adaptive risk assessment frameworks that can address both traditional and technology-driven operational risks (Huang, 2024; Jawdani et al, 2024).

Collectively, these developments signify a paradigm shift: operational risk is no longer a residual or secondary risk category but has emerged as a strategic determinant of institutional resilience, regulatory compliance, competitive positioning, and systemic stability (Chernobai et al., 2018; Curti et al., 2022). This elevation reflects broader recognition that operational failures can precipitate substantial financial losses, reputational damage, regulatory sanctions, and even systemic contagion (Cummins et al., 2006). Within this context, one of the most significant methodological challenges remains the accurate quantification and modeling of operational loss severity distributions, particularly for tail events that are statistically rare yet potentially catastrophic (Dahen & Dionne, 2010).

1.2. Machine Learning and Artificial Intelligence in Operational Risk Modeling

Machine learning (ML) and artificial intelligence (AI) techniques have emerged as a major methodological frontier in operational risk analytics. Their main advantage lies in their ability to model complex, non-linear structures in loss data that are often difficult to capture through traditional parametric approaches. Methods such as Random Forests, Gradient Boosting Machines (GBMs), Support Vector Machines (SVMs), Neural Networks, Deep Learning architectures, and ensemble models offer greater flexibility in identifying interactions and contextual dependencies. They therefore offer a powerful complement to conventional modeling frameworks in the analysis of operational risk. (Drydakis, 2022; Huang, 2024; Pacek et al., 2025; Petropoulos et al., 2020; Sarkar et al., 2016).

Empirical evidence suggests that machine learning can improve the prediction of both operational loss severity and frequency. In particular, Random Forest and Gradient Boosting perform well because they can handle mixed predictors, capture non-linearities, and identify interaction effects. Neural networks may further uncover latent risk patterns that traditional regression models often fail to detect. (Liu et al., 2025; Narang et al., 2025; Zhu et al., 2025).

Recent studies also show that ensemble methods can further improve out-of-sample predictive performance by combining the strengths of multiple base models. In parallel, machine-learning classification algorithms have enhanced the categorization of operational risk events, allowing more granular risk profiling, more targeted mitigation strategies, and lower false-positive rates in automated monitoring systems (Eling & Lehmann,2018; Hassani et al., 2020; Liu et al., 2025; Petropoulos et al., 2020; Rawindaran et al., 2021).

Deep learning has further expanded these possibilities, especially in settings involving sequential dependencies or high-dimensional data. In parallel, natural language processing (NLP) has opened new avenues for extracting risk-relevant information from unstructured sources such as incident reports, audit documents, regulatory texts, and news flows. These developments support more proactive risk management by improving early signal detection and enriching operational risk assessment (Kulkarni & Chandra, 2025; Mienye et al., 2024).

Although ML and AI models offer strong predictive advantages, they also raise important concerns regarding interpretability, transparency, and regulatory acceptance. These issues are particularly relevant in operational risk, where model governance and explainability remain essential. In response, explainable AI (XAI) techniques such as SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), Partial Dependence Plots (PDP), and Individual Conditional Expectation (ICE) have gained importance because they help clarify the contribution of individual variables to model predictions. As a result, XAI frameworks are increasingly seen as a way to reconcile predictive performance with interpretability and regulatory expectations (Ahmed et al., 2022; Alonso & Carbó, 2020; Bodria et al., 2023; Molnar, 2020).

The absence of a universally dominant modeling approach has encouraged the development of hybrid frameworks that combine traditional statistical methods with machine learning techniques. These approaches aim to retain the interpretability and theoretical grounding of classical models while benefiting from the flexibility and predictive capacity of data-driven algorithms. However, important challenges remain. Data scarcity, especially for extreme loss events, may reduce the effectiveness of complex ML models and increase the risk of overfitting. In addition, model validation remains difficult in operational risk settings, particularly under changing business conditions, technological developments, regulatory shifts, and emerging threats. As a result, adaptive and robust modeling strategies continue to represent an important direction for future research. (Hassani et al., 2020; Melina et al., 2025; Petropoulos et al., 2020).

2. Materials and Methods

As illustrated in Figure 1, this chapter adopts a comparative empirical framework to model operational loss severity, drawing on the SAS OpRisk Global Dataset. After filtering and anonymizing the source data, the analysis is restricted to banking-sector loss events recorded between 1996 and 2023. The selected sample provides sufficient temporal and geographic heterogeneity to capture the main empirical features of operational losses, notably strong skewness, heavy tails, and event-level heterogeneity (Moscadelli, 2004).

The entire analysis was implemented using Python, leveraging advanced data science libraries such as Scikit-learn, Pandas, and NumPy to ensure computational efficiency and methodological rigor. The independent variables encompass a comprehensive range of operational risk dimensions, including: Event risk categories, Business lines (Commercial and

Retail Banking), Geographical locations, Firm-level financial indicators, Organizational size and complexity, Regulatory environment characteristics.

The empirical strategy unfolds in three complementary stages. First, a descriptive analysis characterizes the unconditional and conditional distributions of loss amounts across event-risk categories and business lines, highlights outliers, and documents heterogeneity and temporal patterns that may affect model choice and evaluation. Second, a parametric severity-fitting stage benchmarks two classical positive-skew distributions commonly used in operational risk: the Lognormal and the Gamma. These specifications serve as baseline parametric references for strictly positive, asymmetric loss data. Goodness-of-fit is assessed using likelihood-based criteria (AIC/BIC), probability plots, and tail diagnostics to evaluate the model's ability to capture extreme outcomes. Third, a structured-data predictive stage estimates and compares supervised models on the same target variable. The dependent variable is log(1 + loss) to mitigate skewness and stabilize estimation. Evaluated models include:

- Ridge regression (Hoerl & Kennard, 1970) as a regularized linear baseline robust to multicollinearity;
- Decision Tree as an interpretable non-linear baseline;
- Random Forest (Breiman, 2001) to capture non-linearities and interactions via bagged trees;
- Gradient Boosting (Friedman, 2001) to iteratively improve predictive fit through residual correction.

In addition, the chapter introduces a hybrid specification that combines the predictions of the traditional regression model and the Random Forest model. The purpose of this hybrid approach is to test whether integrating linear statistical structure with non-linear machine-learning patterns can improve predictive performance compared with single structured-data models. The dependent variable is the current loss value, and the predictive stage uses its log-transformed form, log(1 + loss), to reduce skewness and improve model stability. Model performance is assessed on hold-out test data using RMSE, MAE, and R-squared. Overall, the methodology seeks to determine whether AI-based structured-data models, whether used alone or in a hybrid form, provide incremental predictive value beyond traditional benchmark approaches.

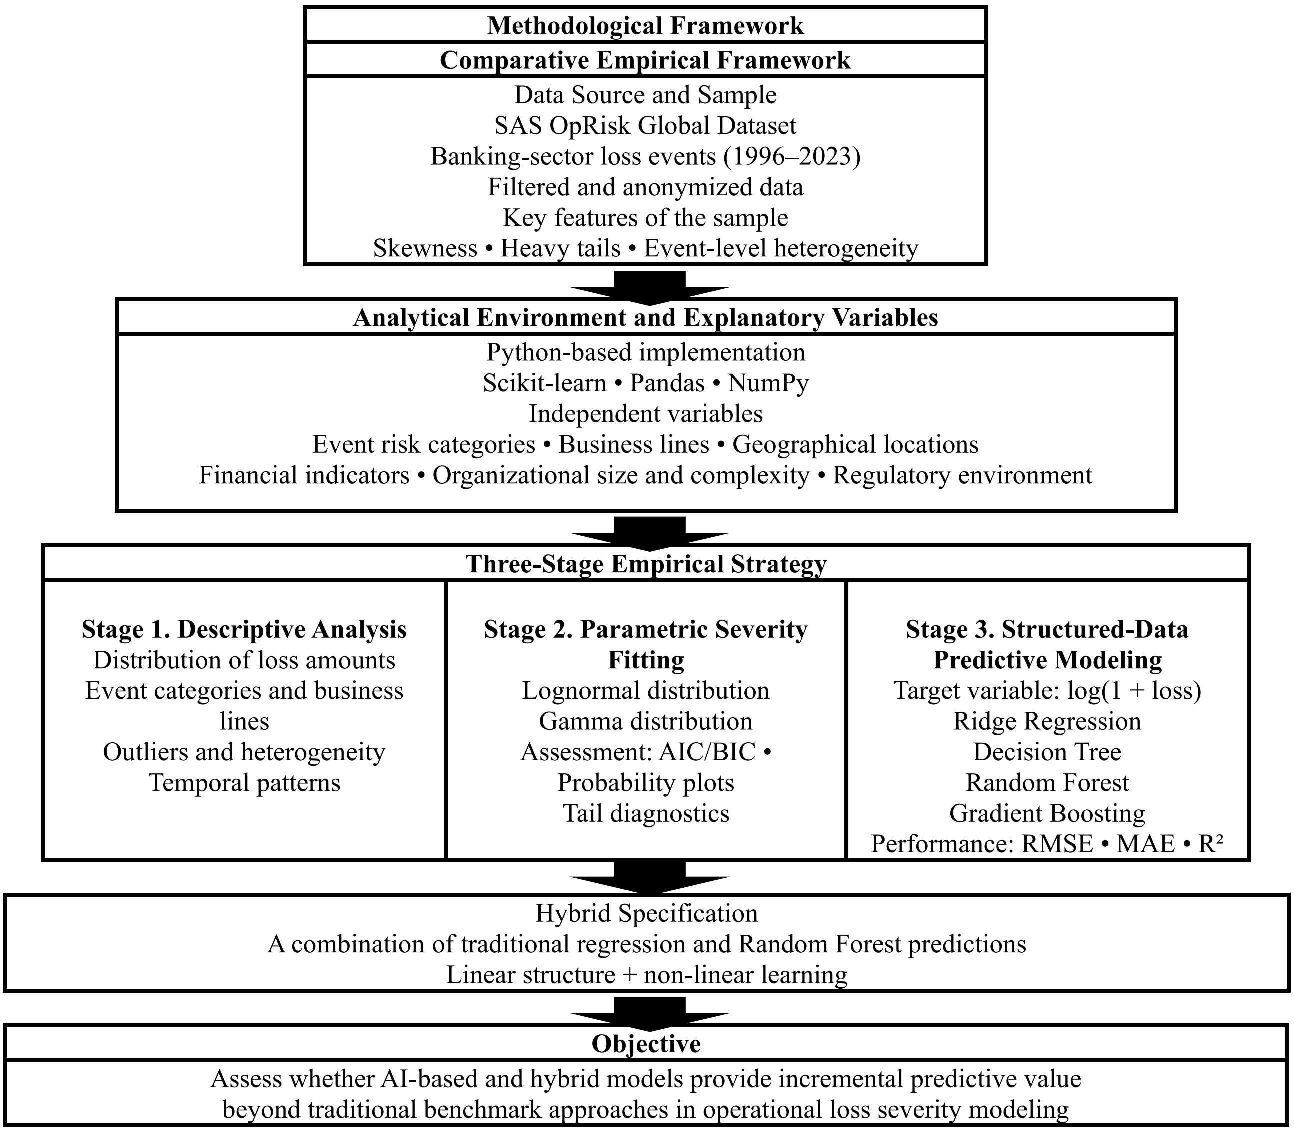

Figure 1: Structure of the Study (Developed by the Authors).

3. Results
3.1. Descriptive Findings
The descriptive analysis provides an initial understanding of the statistical properties of operational loss severity and helps justify the subsequent modeling choices. Before comparing traditional parametric distributions and machine-learning approaches, it is necessary to examine the overall shape of the loss data, its degree of dispersion, and its variation across operational risk categories. The descriptive evidence shows that operational losses are highly asymmetric, highly dispersed, and heterogeneous across event types, making severity modeling particularly challenging.

3.1.1. Distributional Characteristics of Operational Losses
The first important result concerns the overall distribution of operational loss severity (Figure 2: Histogram of Current Value of Loss ($M)). The raw monetary values of loss exhibit pronounced right skew, with a large concentration of relatively small and medium-sized losses and a limited number of extremely large events. This confirms that operational loss severity is not normally distributed and that extreme observations substantially influence the average loss level. The descriptive statistics reinforce this observation. The mean loss is considerably higher than the median, indicating that the distribution is strongly skewed upward by a few large losses. In addition, the gap between the upper quartile and the maximum value highlights the presence of an extended upper tail. These characteristics are typical of operational loss databases and consistent with the heavy-tailed nature of operational risk, as commonly documented in the literature.

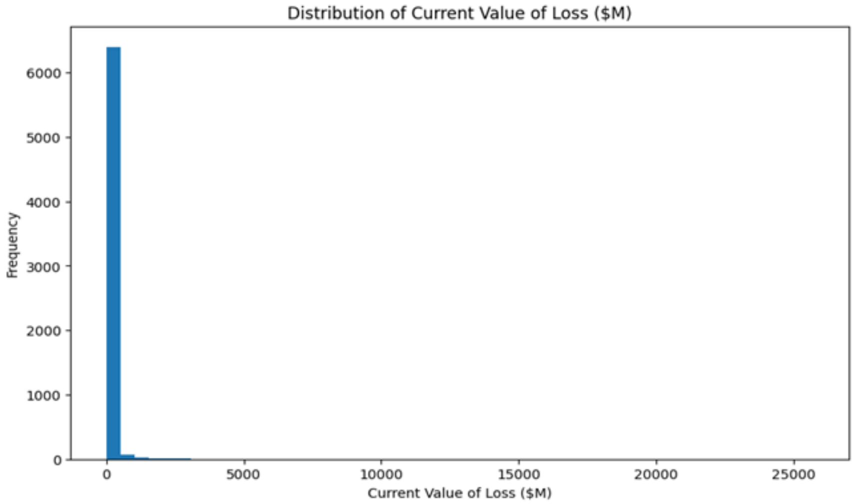

Figure 2: Histogram of Current Value of Loss ($M).

Given this strong asymmetry, a logarithmic transformation of the target variable was introduced for the predictive modeling stage. The log-transformed severity remains asymmetric, as shown in Figure 3, but the transformation substantially reduces dispersion and improves the distribution's regularity. This makes the transformed variable more suitable for supervised prediction models while preserving the fundamental information contained in the loss amounts.

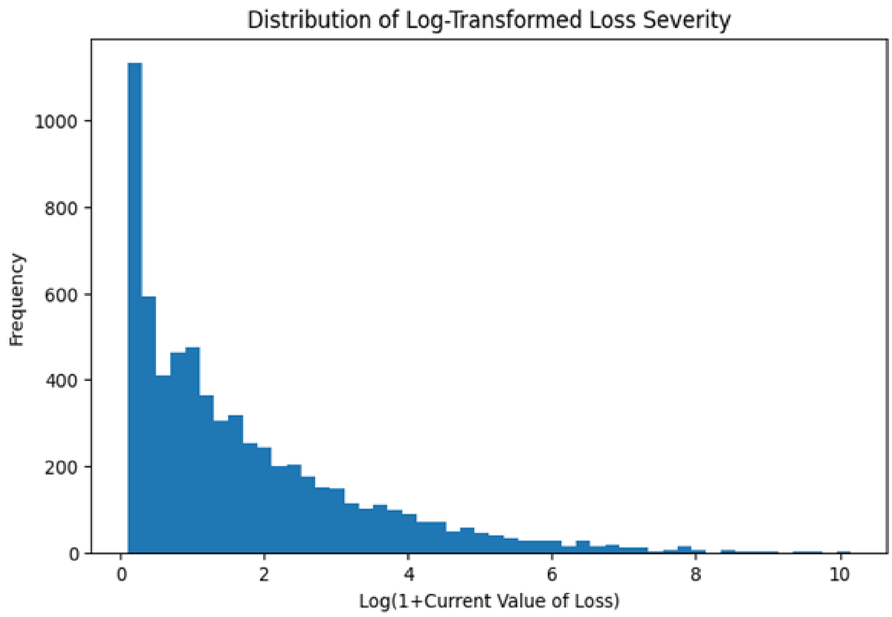

Figure 3: Histogram of Log-Transformed Loss Severity.

3.1.2. Summary Descriptive Statistics

Table 1 summarizes the main descriptive statistics for the loss severity variable. The sample includes 6 535 positive operational loss events. The average loss is markedly higher than the median, indicating that large losses skew the distribution. The standard deviation is much larger than the central quantiles, reflecting the dataset's high dispersion. The minimum observed loss remains very small relative to the maximum, illustrating the sample's wide range.

Table 1: Summary Descriptive Statistics of Operational Loss Severity.

Statistic	Current Value of Loss ($M)	Log-Transformed Severity
Count	6535	6535
Mean	66.453	1.708
Standard Deviation	661.412	1.597
Minimum	0.100	0.095
25th Percentile	0.578	0.456
Median	2.324	1.201
75th Percentile	10.616	2.452
Maximum	25734.045	10.156

3.1.3. Severity Across Event Risk Categories

Operational loss severity also varies substantially across event risk categories. As illustrated in Figure 4, the boxplot of log-transformed severity reveals marked heterogeneity in both central tendency and dispersion across categories. Some categories display higher median loss severity and broader interquartile ranges, whereas others exhibit a lower central tendency but still contain a considerable number of extreme observations. In particular, Clients, Products & Business Practices, and Internal Fraud stand out because of their pronounced upper tails and numerous outliers, indicating a high concentration of large and extreme losses. By contrast, categories such as Execution, Delivery & Process Management, and External Fraud show lower medians, although they remain characterized by a non-negligible degree of dispersion and tail risk.

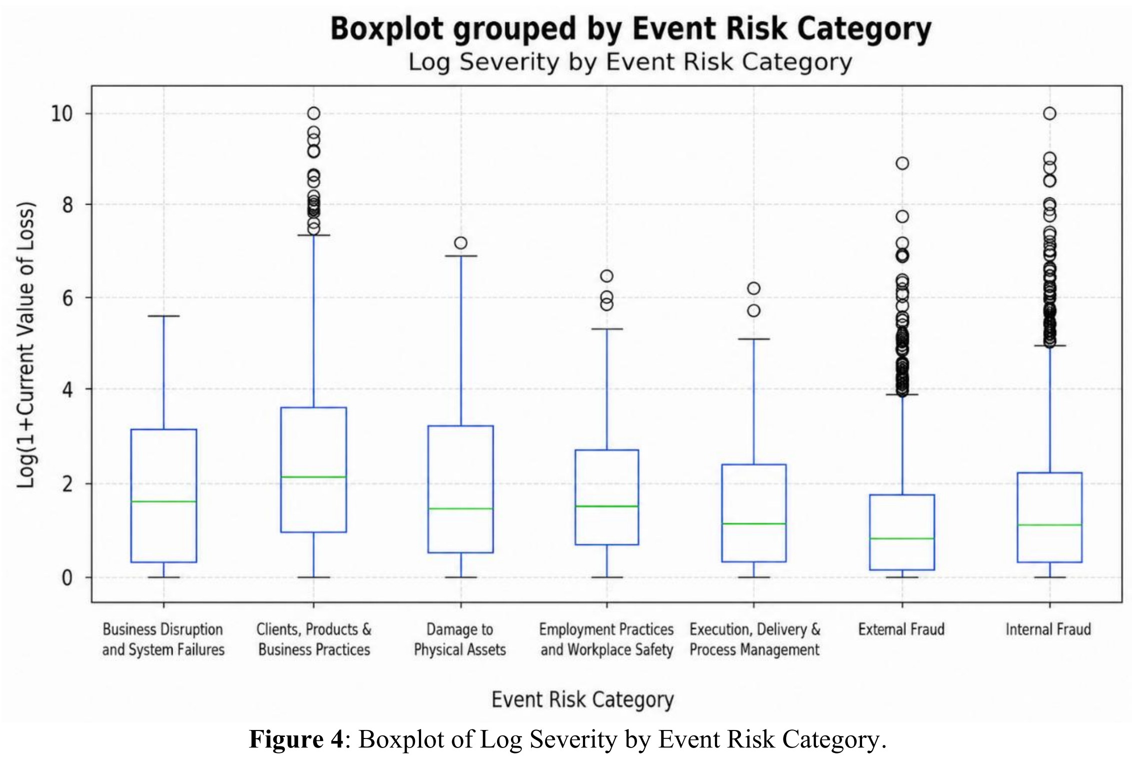

Figure 4: Boxplot of Log Severity by Event Risk Category.

This cross-category heterogeneity is further confirmed by the aggregate severity rankings reported in Figure 5, which presents the event risk categories ordered by mean loss severity. The figure shows that Clients, Products & Business Practices records the highest mean severity, followed by Damage to Physical Assets and Internal Fraud, whereas External Fraud and Execution, Delivery & Process Management are associated with comparatively lower mean loss levels. In parallel, the descriptive statistics shown in Table 1 reinforce these visual patterns, indicating that categories differ not only in mean severity but also in median severity and maximum observed losses. For example, while Internal Fraud combines a relatively high mean severity with very large maximum losses, Clients, Products & Business Practices emerges as the most severe category overall on average.

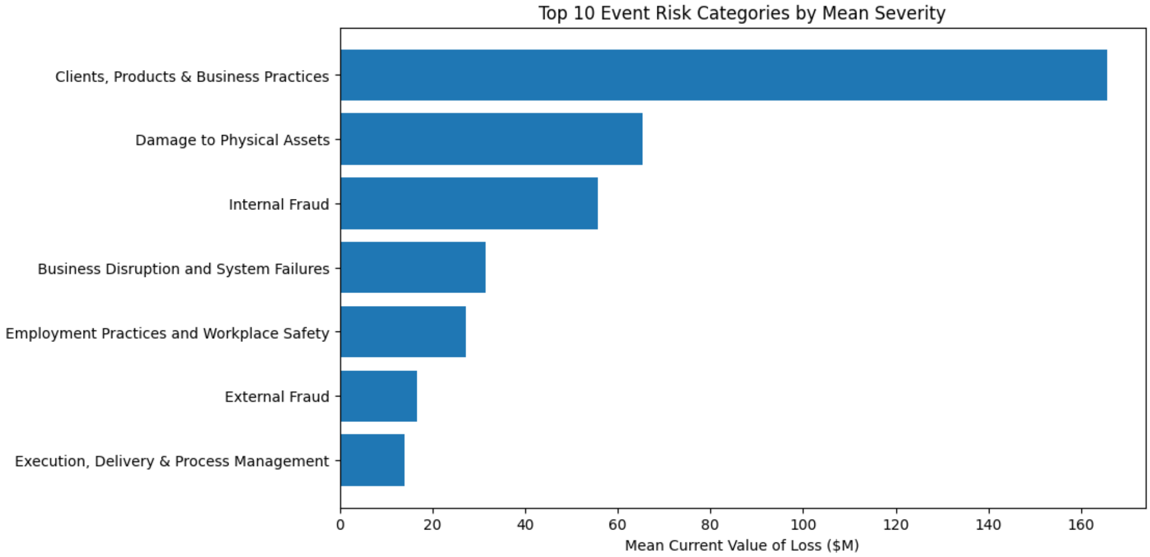

Figure 5: Event Risk Categories Ranked by Mean Loss Severity.

Table 2: Descriptive Statistics of Loss Severity by Event Risk Category.

	count	mean	median	max
Event Risk Category				
Clients, Products & Business Practices	1622	165.591500	7.580141	25734.044680
Damage to Physical Assets	78	65.378852	3.257589	1301.284274
Internal Fraud	2025	55.638506	1.958430	21071.338106
Business Disruption and System Failures	36	31.563626	3.925084	257.487813
Employment Practices and Workplace Safety	110	27.218973	3.456055	626.741364
External Fraud	2506	16.590668	1.444447	6404.636070
Execution, Delivery & Process Management	158	13.947380	2.452374	433.230000

Taken together, Figures 4 and 5 and Table 1 show that operational loss severity cannot be treated as a homogeneous phenomenon across event types. This heterogeneity is important for at least two reasons. First, it confirms that the financial impact of operational events depends not only on their occurrence but also on their underlying nature. Second, it supports the inclusion of event-category information as a central explanatory factor in predictive modeling. The descriptive evidence, therefore, suggests that more flexible modeling approaches may be required in order to capture differences in severity across operational contexts.

3.2. Traditional Severity Distribution Fitting

To complement the descriptive analysis, the study next examines the adequacy of standard parametric distributions for modeling operational loss severity. In operational risk practice, parametric severity distributions remain widely used because they provide interpretable and tractable representations of loss behavior. Among the most common candidates, the Lognormal and Gamma distributions are often used for positive, skewed loss data. The purpose of this section is therefore to assess which of these two traditional specifications offers the better fit for the severity values observed in the sample. The comparison was performed using three complementary criteria: log-likelihood, Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC). A higher log-likelihood indicates a better fit to the data, whereas lower AIC and BIC values indicate a more favorable balance between goodness of fit and model parsimony.

The results clearly show that the Lognormal distribution provides a substantially better fit than the Gamma distribution. As reported in Table 2, the Lognormal model yields a higher log-likelihood and lower AIC and BIC values than the Gamma alternative. This finding indicates that the empirical structure of operational loss severity in the present dataset is more consistent with a Lognormal distribution than with a Gamma distribution.

Table 2: Comparison of Lognormal and Gamma Severity Fits.

Distribution	Log-Likelihood	AIC	BIC
Lognormal	**-21275.330**	**42554.661**	**42568.231**
Gamma	-24483.419	48970.839	48984.408

This result is coherent with the descriptive evidence presented in the previous section. The strong right-skewness of the loss data, the wide dispersion between moderate and extreme losses, and the multiplicative nature of severity variation all point toward a Lognormal specification. By contrast, the Gamma distribution appears less capable of reproducing the observed configuration of the loss data, at least in this sample. At the same time, the comparison should be interpreted within its proper scope. The Lognormal and Gamma models are evaluated here as distribution-fitting tools, not as supervised predictive models. Their role is to describe the general statistical shape of operational loss severity rather than to predict individual losses from explanatory variables. For this reason, they are assessed using likelihood-based criteria instead of predictive performance metrics such as RMSE, MAE, or R2. This distinction is important because a distribution that fits the overall loss profile well does not necessarily outperform more flexible models in supervised prediction tasks.

3.3. Predictive Model Comparison

To move beyond global distribution fitting, the study next compares several supervised predictive models designed to explain and predict operational loss severity from explanatory variables. Unlike the traditional parametric distributions discussed in the previous subsection, these models use structured predictors to estimate the transformed severity of individual operational loss events. This predictive perspective is particularly relevant when the objective is not only to characterize the overall shape of loss severity, but also to identify the modeling strategy that best captures the relationship between loss magnitude and the available explanatory information.

The predictive comparison was conducted using the log-transformed severity variable as the target. This transformation was introduced to reduce the influence of extreme values, stabilize dispersion, and improve comparability across model families. Five structured-data approaches were considered: Ridge regression, Decision Tree, Random Forest, Gradient Boosting, and a model-level hybrid combining Ridge regression and Random Forest. These models were selected because they represent distinct methodological families, ranging from linear statistical benchmarks to nonlinear and ensemble-based machine learning approaches. The hybrid specification was designed to combine the linear component captured by Ridge regression with the non-linear and interaction effects captured by Random Forest. Model performance was evaluated using three standard predictive metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R2). Lower RMSE and MAE values indicate better predictive accuracy, whereas a higher R2 reflects stronger explanatory performance.

The results reported in Table 3 reveal meaningful differences across models. Among the single structured-data models, Random Forest delivers the strongest predictive performance. It achieves lower prediction errors and a higher R2 than Ridge regression, Decision Tree, and Gradient Boosting, indicating that it is better able to capture the complexity of operational loss severity than the other individual structured-data models considered. This result suggests that the underlying relationship between severity and the explanatory variables is not purely linear and likely involves non-linearities and interactions that are better handled by ensemble methods.

Table 3: Predictive Model Comparison.

Model	RMSE	MAE	R^2
Ridge (tabular)	1.355	0.994	0.313
Decision Tree	1.483	1.081	0.177
Random Forest	1.298	0.924	0.369
Gradient Boosting	1.352	0.986	0.316
Hybrid Ridge + Random Forest	**1.273**	**0.922**	**0.393**

Gradient Boosting and Ridge regression achieve relatively similar performance, with only modest differences between them. This indicates that while more flexible algorithms may improve prediction, not all machine-learning models generate equally strong gains in this context. The similarity between these two models also suggests that part of the predictive structure remains accessible to simpler benchmark approaches, even if their overall explanatory power remains lower than that of Random Forest. By contrast, the Decision Tree model records the weakest predictive performance among the structured-data approaches. Although it remains attractive from an interpretability standpoint, its lower R2 and higher prediction errors indicate that a single-tree structure is not sufficiently flexible to capture the full complexity of operational loss severity. This finding points to an important methodological trade-off: the most interpretable model is not necessarily the most accurate one.

Importantly, the best overall structured-data performance is achieved by the hybrid model combining Ridge regression and Random Forest. This hybrid yields the lowest RMSE, MAE, and highest R2 among all structured-data specifications. The improvement over Random Forest alone suggests that linear and non-linear components may contain complementary predictive information. In other words, combining a regularized linear benchmark with a flexible ensemble method appears to enhance predictive performance beyond what either modeling approach can achieve on its own.

4. Discussion

The empirical findings of this study provide nuanced insights into operational loss severity modeling, critically extending and challenging existing literature on risk quantification. Our results demonstrate that machine learning approaches, particularly ensemble methods like Random Forest, substantially outperform traditional parametric distributions in capturing the complex, non-linear dynamics of operational losses. The Random Forest model's R^2 of 0.369 represents a significant improvement over linear benchmarks, confirming previous observations by Chernobai et al. (2011) on the importance of capturing intricate interactions in loss-generating processes.

The hybrid modeling approach, which integrates Ridge regression and Random Forest predictions, emerges as a particularly compelling methodological innovation. Achieving an R^2 of 0.393, this specification not only outperforms individual models but also validates the theoretical arguments proposed by Eling & Lehmann (2018) about the potential of model-level fusion techniques. This finding suggests that operational risk modeling benefits from methodological pluralism, where different modeling logics can provide complementary predictive information.

The stark performance difference between the Decision Tree ($R^2 = 0.177$) and Random Forest ($R^2 = 0.369$) models illuminates a critical methodological trade-off between model interpretability and predictive accuracy. This observation aligns with Sarkar et al. (2018)'s argument that while transparency remains crucial in risk modeling, more flexible

algorithms can capture complexity that simpler, more interpretable models miss. The results underscore the need to move beyond traditional linear approaches while maintaining a nuanced understanding of model limitations.
Importantly, the analysis confirms and extends the descriptive findings from earlier operational risk literature. The pronounced right-skewness and heavy-tailed nature of the loss data, initially evident in the distributional analysis, are methodologically validated by the superior performance of non-linear modeling techniques. The log-transformation of the severity variable emerges as a critical methodological innovation, enabling more stable and meaningful predictive modeling.
The geographical and categorical heterogeneity of operational losses—spanning different business lines, event types, and institutional contexts—further complicates traditional modeling approaches. Our results suggest that machine learning techniques, particularly ensemble methods, can navigate this complexity more effectively, offering a more dynamic and contextually sensitive approach to estimating loss severity. This finding resonates with recent calls in the risk management literature for more adaptive, data-driven modeling strategies.

5. Conclusions

This research has provided a comprehensive examination of operational loss severity modeling, systematically comparing traditional parametric distributions with advanced machine learning approaches, and delivering three principal contributions to the operational risk literature. First, the study provides robust empirical evidence of the pronounced heterogeneity and nonlinearity in operational loss severity, challenging traditional parametric modeling assumptions. Second, by leveraging the extensive SAS OpRisk Global Dataset, the research reveals substantial predictive advantages of machine learning techniques, with the Random Forest model explaining 36.9% of the variance—a significant improvement over traditional linear benchmarks. Third, the study illuminates the critical trade-off between model interpretability and predictive accuracy, demonstrating that while the Decision Tree model offers maximum transparency, its limited predictive performance ($R^2 = 0.177$) emphasizes that the most interpretable models are not necessarily the most accurate. The hybrid modeling approach, integrating linear and nonlinear predictions, underscores the potential of methodological pluralism in risk modeling, suggesting that financial institutions should invest in advanced machine learning capabilities, develop more sophisticated risk quantification strategies, and maintain a balanced approach that combines interpretability with predictive power.
Despite these contributions, the chapter also has limitations. The traditional parametric comparison was restricted to Lognormal and Gamma distributions, and more tail-sensitive approaches may be needed for a deeper treatment of extreme losses. In addition, the predictive analysis focused on structured variables and a limited set of model combinations. Future research could extend this framework in several directions, including more explicit tail modeling, enhanced interpretability tools for machine-learning models, and more advanced hybrid strategies. Further work could also examine the extent to which temporal effects, regulatory changes, and broader contextual factors alter the predictive structure of operational loss severity over time.
By challenging existing operational risk modeling paradigms and providing actionable guidance, the research highlights a fundamental transformation in risk management methodologies—moving from static, parametric distributions to dynamic, data-driven models that capture the complex, ever-evolving nature of operational risks.

Acknowledgements
The authors have no acknowledgments to declare.

Funding
This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Conflicts of Interest
The author(s) report no conflicts of interest.

Data Availability Statement
The data used in this study are not publicly available due to confidentiality requirements and the banking sector's anonymity clause. Operational loss data are highly sensitive and may reveal information about internal processes, control failures, incidents, financial impacts, and banks' risk exposures. Therefore, to preserve the anonymity of the banking institutions concerned, only aggregated and anonymized results are reported

Institutional Review Board Statement
Not applicable.

CRediT Author Statement
Sara CHELH: Conceptualisation; Methodology; Data curation; Formal analysis; Investigation; Resources; Software; Visualisation; Project administration; Writing – original draft.
Mariame ABABOU: Conceptualisation; Methodology; Investigation; Validation; Supervision; Writing – review.
Bouteina EL GHARBAOUI: Conceptualisation; Validation; Supervision; Writing – review & editing.
All authors have read and approved the final version of the manuscript.

Reference List

Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. Research in International Business and Finance, 61, 101646. https://doi.org/10.1016/j.ribaf.2022.101646

Aldasoro, I., Gambacorta, L., Giudici, P., & Leach, T. (2020). Operational and cyber risks in the financial sector.

Alonso, A., & Carbo, J. M. (2020). Machine learning in credit risk: measuring the dilemma between prediction and supervisory cost.http://dx.doi.org/10.2139/ssrn.3724374

Basel Committee on Banking Supervision. (2021). *Revisions to the principles for the sound management of operational risk*. Bank for International Settlements

Basel Committee on Banking Supervision. (2024). *Core principles for effective banking supervision*. Bank for International Settlements.

Board, F. S. (2020). Effective practices for cyber incident response and recovery: Final Report. Basel, October 19.

Bodria, F., Giannotti, F., Guidotti, R., Naretto, F., Pedreschi, D., & Rinzivillo, S. (2023). Benchmarking and survey of explanation methods for black box models: F. Bodria et al. Data Mining and Knowledge Discovery, 37(5), 1719-1778. https://doi.org/10.1007/s10618-023-00933-9

Chernobai, A., Jorion, P., & Yu, F. (2011). The determinants of operational risk in US financial institutions. Journal of Financial and Quantitative Analysis, 46(6), 1683-1725.https://doi.org/10.1017/S0022109011000500

Cope, E. W., Piche, M. T., & Walter, J. S. (2012). Macroenvironmental determinants of operational loss severity. Journal of Banking & Finance, 36(5), 1362-1380.https://doi.org/10.1016/j.jbankfin.2011.11.022

Cruz, M. G., Peters, G. W., & Shevchenko, P. V. (2015). Fundamental aspects of operational risk and insurance analytics: A handbook of operational risk. John Wiley & Sons.

Cummins, J. D., Lewis, C. M., & Wei, R. (2006). The market value impact of operational loss events for US banks and insurers. Journal of Banking & Finance, 30(10), 2605-2634. https://doi.org/10.1016/j.jbankfin.2005.09.015

Curti, F., Frame, W. S., & Mihov, A. (2022). Are the largest banking organizations operationally more risky?. *Journal of Money, Credit and Banking*, *54*(5), 1223-1259.https://doi.org/10.1111/jmcb.12933

Dahen, H., & Dionne, G. (2010). Scaling models for the severity and frequency of external operational loss data. Journal of Banking & Finance, 34(7), 1484-1496.https://doi.org/10.1016/j.jbankfin.2009.08.017

Drydakis, N. (2022). Artificial intelligence and reduced SMEs' business risks. A dynamic capabilities analysis during the COVID-19 pandemic. Information Systems Frontiers, 24(3), 1223-1247. https://doi.org/10.1007/s10796-022-10249-6

Dutta, K., & Perry, J. (2006). A tale of tails: an empirical analysis of loss distribution models for estimating operational risk capital (No. 06-13). Working papers.

Eling, M., & Lehmann, M. (2018). The impact of digitalization on the insurance value chain and the insurability of risks. The Geneva Papers on Risk and Insurance, 43(3), 359-396.https://doi.org/10.1057/s41288-017-0073-0

Financial Stability Board. (2023). *Enhancing third-party risk management and oversight: A toolkit for financial institutions and financial authorities*

Hassani, H., Huang, X., Silva, E., & Ghodsi, M. (2020). Deep learning and implementations in banking. Annals of Data Science, 7(3), 433-446. https://doi.org/10.1007/s40745-020-00300-1

Huang, X. (2024). Model Risk Management in the Era of Artificial Intelligence (Master's thesis, Universidade NOVA de Lisboa (Portugal)).

Jadwani, B., Parkhi, S., & Mitra, P. K. (2024). Operational risk management in banks: A bibliometric analysis and opportunities for future research. *Journal of Risk and Financial Management*, *17*(3), 95. https://doi.org/10.3390/jrfm17030095

Kulkarni, O., & Chandra, R. (2025). DynBERG: Dynamic BERT-based Graph neural network for financial fraud detection. arXiv preprint arXiv:2511.00047. https://doi.org/10.48550/arXiv.2511.00047

Liu, S., & Wang, C. (2025). Research on External Risk Prediction of Belt and Road Initiative Major Projects Based on Machine Learning. Sustainability, 17(20), 9089. https://doi.org/10.3390/su17209089

Liu, C., Chen, X., Wang, Y., Xiao, S., Liu, C., & Liu, F. (2025). Exploring the determinants of systematic risk through machine learning: evidence from Chinese listed companies. Applied Economics, 57(59), 10454-10468.https://doi.org/10.1080/00036846.2024.2431170

Melina, M., Sukono, Napitupulu, H., & Mohamed, N. (2025). Modeling of Machine Learning-Based Extreme Value Theory in Stock Investment Risk Prediction: A Systematic Literature Review. Big Data, 13(3), 161-180.https://doi.org/10.1089/big.2023.0004

Molnar, C. (2020). Interpretable machine learning. Lulu. com.

Moscadelli, M. (2004). The modelling of operational risk: experience with the analysis of the data collected by the Basel Committee. Available at SSRN 557214. http://dx.doi.org/10.2139/ssrn.557214

Mienye, E., Jere, N., Obaido, G., Mienye, I. D., & Aruleba, K. (2024). Deep learning in finance: A survey of applications and techniques. Ai, 5(4), 2066-2091. https://doi.org/10.3390/ai5040101

Moosa, I. A. (2007). Operational risk: a survey. Financial markets, institutions & instruments, 16(4), 167-200. https://doi.org/10.1111/j.1468-0416.2007.00123

Narang, R. (2025). Comparative Analysis of AI-Based Driving Behavior Risk Assessment Methods Accuracy, Efficiency, and Real-World Applicability. Spectrum of Research, 5(2).

Pacek, Ł., Mróz-Gorgoń, B., & Czugan, M. (2025). Risk management in an organization in the era of AI–a theoretical approach. Zeszyty Naukowe Politechniki Poznańskiej. Organizacja i Zarządzanie.

Petropoulos, A., Siakoulis, V., Stavroulakis, E., & Vlachogiannakis, N. E. (2020). Predicting bank insolvencies using machine learning techniques. International Journal of Forecasting, 36(3), 1092-https://doi.org/1113.10.1016/j.ijforecast.2019.11.005

Rawindaran, N., Jayal, A., & Prakash, E. (2021). Machine learning cybersecurity adoption in small and medium enterprises in developed countries. Computers, 10(11), https://doi.org/150. 10.3390/computers10110150

Sarkar, S., Vinay, S., Raj, R., Maiti, J., & Mitra, P. (2016). Application of optimized machine learning techniques for prediction of occupational accidents. Computers & Operations Research, 106, 210-224. https://doi.org/10.1016/j.cor.2018.02.021

Shehadeh, M., Al-Nadabi, A. K., Yaseen, H., & Migdady, H. (2026). AI-based ensemble learning model for financial leverage prediction and risk management in thebanking sector. Journal of Financial Reporting and Accounting, 1-23. https://doi.org/10.1108/JFRA-01-2025-0013

Zhu, Z., Tong, P., Dai, Z., Zheng, M., Tang, W., Zhu, T., & Ding, R. (2025, August). Development of Operational Risk Warning of Fixed Assets Based on Improved Machine Learning Model. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-5). IEEE. https://doi.org/10.1109/IACIS65746.2025.11211348