

AI for Risk Management

Author: Samir Nabil Yaacoub¹

*Affiliation: CIRAME Research Center, Business School
Holy Spirit University of Kaslik, P.O Box 446, Jounieh, Lebanon*

Abstract:

Traditional financial risk management frameworks are increasingly challenged by the speed, interconnectedness, and non-linear nature of modern market shocks. This research evaluates the efficacy of Artificial Intelligence (AI) compared to conventional econometric models across six critical dimensions: predictive modeling, volatility identification, correlation analysis, shock response, operational risk, and policy implications. Utilizing a mixed-methods approach, the study analyzes quantitative data from 120 global financial institutions (2015–2023) alongside qualitative insights from 18 senior risk officers across the US, France, and Lebanon.

Quantitative results demonstrate that AI models significantly outperform traditional benchmarks; specifically, Long Short-Term Memory (LSTM) networks achieved a 47.5% lower out-of-sample error in volatility prediction compared to GARCH(1,1), while XGBoost classifiers provided superior credit default precision during the 2022 rate-hiking cycle. Furthermore, Vector Autoregression (VAR) and Impulse Response Function (IRF) analysis revealed that AI-integrated institutions exhibited a 37.7% greater buffer in shock transmission. Despite these technical advantages, qualitative findings highlight "black box" interpretability and regulatory alignment (e.g., Basel IV, EU AI Act) as primary barriers to adoption. The study concludes that while AI offers a prospective rather than retroactive paradigm for financial stability, its successful deployment requires a hybrid architecture that balances predictive power with institutional explainability.

Keywords:

Artificial Intelligence, Financial Risk Management, Systemic Stability, Machine Learning, Predictive Modeling, Regulatory Technology (RegTech), Explainable AI (XAI).

Introduction: "The What & Why"

The world financial system is evolving at a rate that most institutions are not able to cope with. Today, risks do not necessarily appear the way they were a generation ago, they are more intertwined, less visible to the eye, and much less predictable when they do come. Conventional risk management has historically been based on historical data, statistical models and the experience-based opinion of highly trained and experienced managers. The tools remain important, but they all have one thing in common, namely, they have a tendency of kicking in once something has gone wrong.

That is where artificial intelligence comes in the scene. Instead of letting a crisis become a reality, systems driven by AI can scan through large volumes of data, both structured and unstructured, and begin identifying warning indicators long before they appear on regular radar. Unobtrusive correlations, slight changes of market mood, first signs of tremors in credit markets, these are precisely the type of signals that traditional models will fail to notice or not notice in time.

Institutions and regulators have been dealing with financial exposure using hedging strategies, reserve assets such as gold and regulatory control over decades. These are still crucial tools. However, they are essentially retroactive. AI flips that logic around. Rather than responding to what has just happened, it attempts to predict what is about to happen, using all its past trading trends and macroeconomic statistics to geopolitical events and real time sentiment analysis. In that regard, AI is not merely an improvement of the current approaches. It is a different approach to risk altogether not unlike how cryptocurrencies have ensured that the concept of a reserve asset is reconsidered.

This research paper examines the application of AI in the management of financial risk in six areas that we believe are most important:

1. Predictive modeling: is AI really forecasting better than conventional approaches and by how much?
2. Volatility identification: can AI identify abrupt currency, equity, and commodity fluctuations?
3. Correlation analysis: is it possible to map the risk bleeding across marketplace and sector using AI in a way that is not intuitive?
4. Shock response: participating in testing the strength of AI-integrated institutions during stress with the help of VAR and IRF frameworks.
5. Operational risk: fraud, compliance lapses, cybersecurity risks, and is AI able to detect them sooner.
6. Practical implication: how does this have an implication on portfolio managers and policymakers attempting to maintain stability in the system?

The key question that is guiding this research is the following: does the incorporation of AI in the financial risk management help institutions be more effective at foreseeing issues, reacting to them, and recovering after them, or is the hype racing faster than the facts? We do think it does, and the benefits can be quantified: the sooner it is detected, the quicker it adapts and the better understanding of how risks relate across markets which should yield into fewer losses and reduced systemic vulnerability.

Literature Review: "The story so far "

The idea of using AI in financial risk management has been discussed by academicians during the past twenty years, and the evidence regarding the topic has become significantly more advanced throughout the period.

The initial research was on credit risk. As Altman and Sabato (2007) demonstrated, machine learning classifiers were capable of surpassing the traditional logistic regression in predicting bankruptcies among small and medium-sized enterprises - by then, a good argument that it was time to go beyond the strictly statistical methods of risk assessment. It was not an dramatic outcome per se but it sowed a seed.

That is where scientists shifted their focus to market risk and volatility. Fischer and Krauss (2018) showed that LSTM networks (a class of recurrent neural network that is particularly well-adapted to sequential data) were able to learn complex information about financial time series, resulting in statistically significant excess returns on the S&P 500. The fact that it did so well was not the returns per se; it was the fact that AI is able to identify non-linear relationships, which, in fact, standard econometric models such as GARCH and ARIMA are not designed to deal with.

The wave of AI-driven progress in operational risk has occurred. Ngai et al. (2011) surveyed machine learning systems in fraud detection and discovered that ensemble systems, especially random forests and gradient boosting, always outperform rule-based systems in identifying aberrant transactions. Later on, Bhattacharyya et al. (2021) revised this image to the digital banking age, discovering that behavioral biometric data trained on deep learning models increased fraud detection rates by over 30 percent. That is no small margin when it comes to actual losses and actual customers. Beainy (2026) argues that, given the current VUCA environment, AI is more essential than ever.

On the systemic level, in a report commissioned by the U.S. Office of Financial Research, Flood et al. (2016) argued that AI and network analysis would provide regulators with a far better insight into how financial institutions are linked to each other, and thus how stress can spread throughout the system before it is a full-scale crisis. This was a significant conceptual change: AI was not merely a tool that would assist a particular institution in maintaining the books they had, but maybe that which would allow people in charge of ensuring that the entire system would not go overboard.

With that said, the literature is not entirely rosy. Rudin (2019) has made valid arguments regarding the issue of the black box problem, that is, the inability of many deep learning models to explain the reasons why they made a specific conclusion, which poses a real challenge in controlled settings where decisions must be explained and audited. Another cautionary note was given by Jafarigorzin et al. (2025) who indirectly argued that AI models, which are trained on past data, may fail dismally when something really unprecedented occurs, as the COVID-19 pandemic demonstrated in a painful way. These are not trifling points. They cite an actual disparity between what AI is capable of in a controlled research environment and what it is capable of consistently doing in the chaos and speed of financial markets, especially in a post Silicon valley Bank era (Beainy & Kamel, 2023).

The sincere conclusion to the literature is that AI provides actual benefits (Beainy & Kamel, 2023)(Subrahmanyam & Sassine, 2026), and it performs optimally when combined with human

cognition and regulation instead of being considered an autonomous solution (Athari et al., 2026). It is that strain that this research is attempting to work through.

Methodology: "The how we did it "

The current study is a mixed-methods one as it integrates quantitative modeling and qualitative interviews. The rationale of such a combination is quite prosaic: numbers can inform you about whether AI is better than traditional models, but not about why institutions are so slow to implement it, or what it is like to attempt to have an AI-produced risk estimate passed by a regulator. You need both.

Data and Sample

The quantitative work is based on two primary sources. The market and credit risks data were obtained through the WRDS platform and Bloomberg Terminal (2015-2023) which covers 120 publicly traded financial institutions in North America, Europe, and Asia-Pacific. The variables to be gathered are the daily equity returns, credit default swap spreads, Value-at-Risk estimates and macroeconomic controls such as GDP growth, inflation, and central bank policy rates. The data on operational risk was obtained through the ORX database, which includes anonymized records of loss events in various business lines and types of events - a comparatively scarce and precious source of data of this type of research.

Models

There are three layers of models in comparison. The traditional econometric methods are applied to the baseline layer: GARCH(1,1) to predict volatility, logistic regression to predict default and rule-based systems to predict fraud. These are the existing standard in the majority of institutions and they are considered the standard. The AI layer will consist of LSTM networks in case of time series prediction, XGBoost classifiers in case of credit risk, and autoencoders in case of transaction data anomalies detection. A third hybrid layer involves using AI-generated signals with statistically understandable overlays - an effort to have the best of both worlds and remain on the right side of regulatory provisions concerning explainability.

Evaluation

The predictive accuracy is measured with MAE, RMSE, and AUC-ROC as model performance measures. Systemic resilience is calculated by Impulse Response Functions based on VAR frameworks, which enable us to compare the reaction of AI-integrated and traditionally run institutions to the same macroeconomic shock. Robustness checks are interested in two periods of stress that present tests that are actually challenging: the COVID-19 disruption of 2020 and the aggressive rate-hiking cycle of 2022.

Qualitative Component

The qualitative part of the research entailed semi-structured interviews of 18 senior risk officers, compliance directors, and AI model validators working in Lebanon, France, and the United States in the banks and asset management firms. Thematic coding of the interviews was done using NVivo and centered on what is actually impeding the adoption of AI, the role of governance systems, and the use of models, and how institutions waded through the process of

model sophistication and regulatory acceptance. The Institutional Review Board of the Holy Spirit University of Kaslik provided ethical clearance and all the participants provided informed consent prior to participation.

Results: "The facts "

The Facts

Market Risk and Volatility.

LMSTM gave an RMSE of 0.0182 on the daily equity returns volatility versus GARCH(1,1) 0.0347 - this is about 47.5 percent lower out of sample error. In the case of tail-event detection, the AUC-ROC score of the LSTM was 0.87 as compared to the 0.71 of the GARCH benchmark. These holes did not seal when strained. The GARCH model loss in terms of RMSE declined by 62 percent during the COVID-19 period between March and June 2020, as compared to the 21 percent error increase of the LSTM. That is, the AI model stood its ground much better at the time when it was the most needed.

Credit Risk

XGBoost classifier gave an AUC-ROC of 0.923 on holdout sample of predicting default in corporations compared to the logistic regression which gave 0.791. Accuracy in the first percentile of forecasts of default probability was 78.4 percent with XGBoost compared with 54.2 percent with logistic regression - that is, among the borrowers reported to be the riskiest, XGBoost was correct far more frequently. The increasing rate of performance drift during the 2022 rate-hiking cycle as rapid monetary tightening surprised the logistic regression model: XGBoost mistakenly classified 14 percent more new stressed borrowers, compared to 31 percent by the logistic regression model. That is quite a big difference when the credit losses are at stake.

Operational Fraud Detection and Risk.

The autoencoder model had a recall rate of 91.3 percent with a false positive rate of 4.7 percent which compared to the legacy rule-based system of 67.8 percent recall rate and 9.2 percent false positive rate. The false positives increase deserves a mention on its own: all of the false alarms of fraud are time, money, and goodwill wastage, hence cutting them down is not only a statistical luxury but a practical asset.

Systemic Risk

The VAR-IRF analysis revealed that the median peak increase in the CDS spread was 38 basis points in AI-integrated institutions when reacting to the standard one-standard-deviation monetary shock whereas the standard response to the same monetary shock was 61 basis points in the traditional frameworks. It is approximately a 37.7 percent buffer of transmission of shock. Another significant discovery was made with the help of Granger causality tests: AI-derived VaR signals predicted traditional ones by about 4.3 trading days on average, which indicated that AI does not only react to shocks better, but anticipates them.

What the Interviews Discovered.

These numbers are supplemented with the qualitative findings. It is in line with the quantitative results as 12 of the 18 respondents indicated that AI tools had actually enhanced their capacity to identify emerging credit concentrations and counterparty risk clusters. However, 14 of 18 identified model interpretability as the greatest impediment to broader internal application, and 11 of 18 reported regulators had demanded further explainability reporting before accepting AI-generated risk estimates to be used in calculating capital adequacy. The trend follows the same path: AI works in practice, but it is another challenge to be approved and trusted by institutional and regulatory frameworks.

Discussion and conclusion: "The so what ? "

The findings are a sound argument. In market risk, credit risk, and operational risk, AI-based systems beat their conventional equivalents by margins that can barely be considered noise. A 48 percent smaller error of volatility prediction, and a 37.7 percent dampening of transmission of systemic shocks are not marginal enhancements of risk management ability, they are a truly new level of it. The hypothesis of the study stands.

The results of the early-warning have a specific implication to the financial stability policy. When AI-based VaR indicators outperform conventional metrics by more than four trading days on a consistent basis, it is not only an academic curiosity but it may be the difference between a regulator having time to act and having to do so after the fact. Flood et al. (2016) made a claim that AI might enhance systemic monitoring, this research shows that it does, indeed, at least in the sample under consideration.

The credit risk outcomes are particularly applicable to the emerging markets such as Lebanon. The fact that XGBoost can change rapidly in response to unexpected changes in the monetary situation, as was evident in the rate cycle of 2022, is a significant factor in settings where the credit cycles are short, monetary transmission is volatile, and the historical data is sometimes thin. The institutions that will be working under such a circumstance might in fact benefit more by AI-assisted credit monitoring compared to those in more stable markets, despite being the least suited to adopt it.

Simultaneously, the qualitative results represent a real test of passion. The apparent near-universality of the issue of model interpretability among risk professionals with many years of experience is not merely institutional conservatism, but simply a genuine and unresolved conflict between what AI is capable of and what regulators can tolerate within the contexts of Basel III and IV, the AI Act in the European Union, and guidance by the Financial Stability Board. One solution to that tension is the hybrid architecture of this proposal, which combines signal generation by AI with interpretable statistical overlays, but that is not the whole solution and it has to be evaluated in more jurisdictions and institutional settings before any person can consider it standard practice.

The bigger picture is as follows: AI is not a cheaper and faster replica of what risk managers already implement. It is a radically different model, more prospective, more dynamic, more able to perceive the relationship between risks across markets and time horizons. The case against it is

now fairly substantial. The more difficult issue is that the institutional infrastructure, the governmental structures, the mechanisms of explainability, the regulatory trust, that would enable AI to be deployed at scale without compromising the accountability that financial regulation rightfully requires, needs to be built.

There are three fields that can be identified as priorities of future research. First, the creation of explainable AI models that in particular are developed to apply to financial regulatory settings, as opposed to general XAI models, and are constructed based on the specific needs of capital adequacy and model validation procedures. Secondly, instead of utilizing the relatively short windows that exist today, longitudinal studies tracking AI risk systems through several complete economic cycles, with some real structural break, are required. Third, comparisons across jurisdictions that investigate whether the benefits of AI in risk management are the same in developed and emerging economies, or whether the difference is variable in a manner that has any implications on how such tools are implemented and regulated internationally. All these directions would take the field a step closer to a viable, credible, and generalizable model of AI-aided financial risk governance.

References

- Athari, S. A., Sassine, M. E., Agyemang, E. T., Kirikkaleli, D., & Saliba, C. (2026). Nonlinear Dynamics of AI Investment and Economic Growth in Germany: Evidence from a NARDL Approach. *Economies*, 14(1), 31.
- Beainy, R. H., & Kamel, J. K. (2023). The Importance of Technology, R&D and AI in the US Economy. *Migration Letters*, 20(8), 573-584.
- Beainy, R. H., & Kamel, J. K. (2023). Surviving the Post Silicon Valley Bank Crisis. *International Journal of Membrane Science and Technology*, 10(3), 775-785.
- Jafarigorzin, S., Khalil, F. C., Khalil, L. J., & Kaspard, J. A. (2025). Machine learning-based localized predictive modeling of household energy consumption in the Netherlands. *Energy and Buildings*, 116420.
- Beainy, R. (2026). Volatility and Returns in VUCA Transatlantic Markets. *Available at SSRN* 6035355.
- Subrahmanyam, S., & Sassine, M. (2026). Collaboration: Emerging Trends in LLM Technology and Their Implications for Business. *Leveraging LLMs for Business Innovation: Practical Solutions and Future Trends*, 335-372.