AI for Financial Stability:
Dynamic Models that Learn and Warn

Paper written by:
Anthony **Carbonelli**
Mario **Di Salvo**

HEPHAESTUS APPLIED ARTIFICIAL INTELLIGENCE ASSOCIATION

Beyond Static Risk Models

Financial risk models were built on the belief that markets behave in stable, predictable patterns. Value-at-Risk and Monte Carlo simulations work by replaying history thousands of times under slightly different assumptions, hoping that the future will behave like a rearranged version of the past. Yet the last two decades have repeatedly shown how fragile that assumption is. The 2008 credit crisis, the rapid market reversals during the pandemic, and the shocks from geopolitical tension all revealed the same flaw: these systems react only after instability has already taken hold.

The consequence is not just technical failure but human cost. When models underestimate risk, banks freeze lending, jobs vanish, and household wealth erodes. Adaptive systems that learn from changing conditions can help institutions respond faster and protect the real economy that depends on them.

Recent advances in artificial intelligence allow risk management to evolve beyond fixed formulas. AI can learn from new information as it arrives, detecting relationships that shift over time and signaling instability before it spreads. The Bank for International Settlements (BIS) demonstrates this potential in its 2025 working paper: "Predicting Financial Market Stress with Machine Learning". There, economists show that machine learning can capture the complex, non-linear connections between liquidity, leverage, and sentiment that older approaches could not.

Dynamic Stress Testing with AI

Stress testing once meant feeding a handful of extreme scenarios, such as a sudden rate hike or a recession, into a model and checking how portfolios would react. These exercises, once designed, could not evolve as markets changed.

AI transforms this process. Machine-learning algorithms can simulate thousands of potential crises instead of just a few, generating an entire distribution of plausible outcomes. They detect interactions between assets, institutions, and behaviors that traditional methods miss: for instance, how margin calls in one market could trigger liquidity shortages elsewhere, creating a living map of potential stress across markets.

The BIS 2025 study demonstrated this evolution by constructing Market Condition Indicators (MCIs) for the Treasury, foreign exchange, and money markets. These indicators measure how smoothly markets are functioning by combining data on volatility, liquidity, and arbitrage efficiency. When dealers stop quoting prices quickly or interest-rate spreads widen abnormally, the MCIs rise, signaling that funding is tightening and liquidity is deteriorating. When a model trained on these indicators predicts future stress, it effectively functions as an early warning system.

To forecast these indicators, researchers used a random forest, a type of machine-learning model that builds many decision trees and averages their predictions. Each tree learns simple if-then rules from historical data, for example, if bond-market liquidity falls and volatility rises, stress increases. By combining hundreds of such trees trained on different data samples, the random forest captures non-linear relationships while avoiding the over-confidence of a single model.

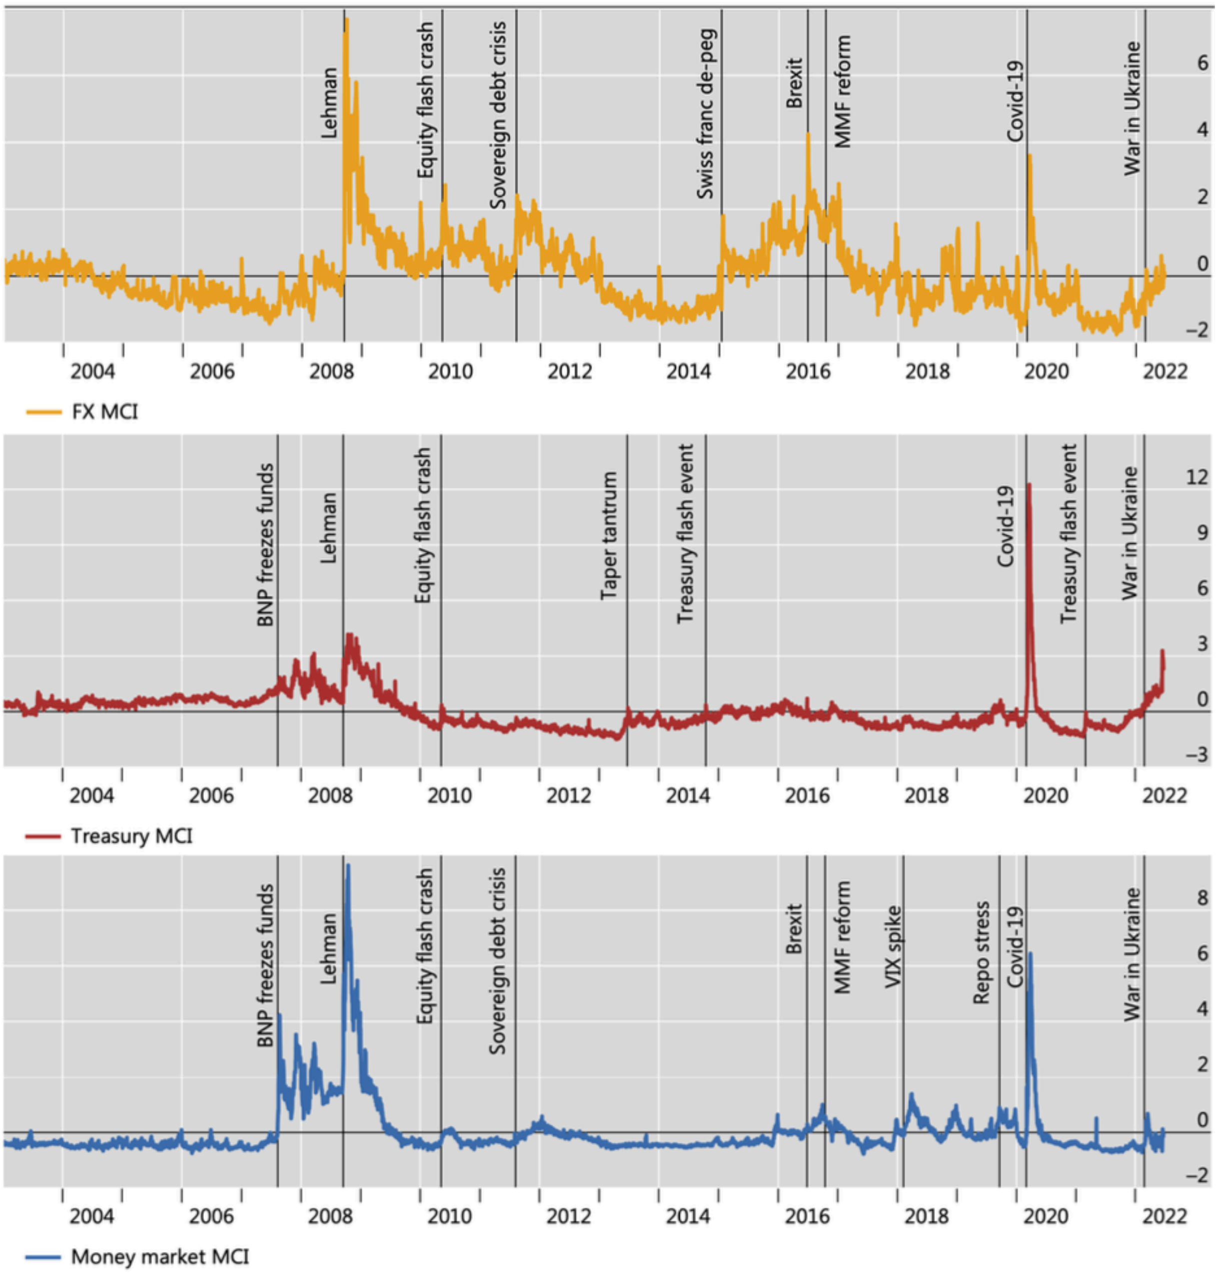

Figure 1: Market Condition Indicators for Treasury, FX, and money markets; higher values indicate tighter conditions. Source: Bank for International Settlements, BIS Quarterly Review (Sep 2022), Graphs 2–4. © BIS.

Spotting Market Stress Early

Financial markets move faster than human perception. Prices change every second, and millions of variables interact in ways that are too complex for traditional analytics. Machine learning models can process vast amounts of time-series data, identifying patterns that precede turbulence.

The BIS research found that investor overextension and funding liquidity constraints consistently predict future market stress. By learning these relationships, AI helps regulators and financial institutions detect early signs of fragility: when liquidity begins to evaporate or when volatility rises independently of macroeconomic news. This allows for preventive action: temporary liquidity injections, position limits, or hedging strategies before instability spreads through the system.

AI's strength lies not in predicting a single crisis but in revealing where and how stress is forming, giving institutions the time to prepare rather than react.

Real-Time Credit Signals

Traditional credit models treat borrowers as fixed points in time, scoring them on past income, leverage, and payment history. Such static profiles fail to capture how behavior changes under stress. AI systems, by contrast, can learn from continuous streams of data: payment patterns, cash-flow irregularities, and market indicators that signal deteriorating creditworthiness before default occurs.

A 2024 Journal of Financial Markets study showed that deep-learning models outperformed conventional credit scoring by roughly 30 percent in predictive accuracy, particularly during downturns. These models interpret credit risk as a dynamic probability that evolves with new information rather than a single rating.

For lenders, this means earlier interventions and fewer unexpected losses. For regulators, it offers a clearer view of systemic exposure. And for borrowers, it can mean fairer credit assessments that reflect real-time behavior rather than outdated history.

Operational Resilience and Systemic Risk

Modern finance depends on complex digital networks. Transactions, settlements, and compliance checks all occur through systems that must function without interruption. AI enhances their resilience by acting as a continuous monitoring layer.

Machine-learning models can process billions of transactions, learning what "normal" looks like and flagging anomalies that may indicate cyber threats, data corruption, or fraud. Some algorithms use reinforcement learning to adapt their detection strategies as attack patterns change. The systems identify operational irregularities in real time and can even trigger automated containment protocols when risk thresholds are breached.

Beyond single institutions, AI also helps identify systemic vulnerabilities. By analyzing correlations between institutions' exposures, liquidity positions, and collateral movements, algorithms can uncover hidden dependencies, such as how stress in money markets might spill over to foreign exchange or government bonds. This interconnected view, once impossible with static tools, is vital for preventing cascading failures in global finance.

Limits and Risks of Algorithmic Forecasting

Despite its capabilities, AI is far from infallible. Markets are driven by human expectations, emotion, and collective behavior, forces that rarely follow patterns visible in data. Algorithms trained solely on the past risk missing turning points born from psychology rather than statistics.

Three challenges stand out:

1. Data bias and incompleteness: Historical records reflect past inequalities and market regimes. Models that learn from them may embed outdated assumptions about risk and behavior.

2. Opacity and explainability: The reasoning of complex models remains statistical, not causal. In risk management, where understanding cause and effect is essential, this opacity can lead to misplaced confidence.
3. Dependence and feedback: The more institutions use similar AI models, the greater the chance that their decisions converge. If many investors rely on the same risk signals, they might react simultaneously, unintentionally amplifying volatility.

Equally important is the role of human judgment. Financial decisions involve ethical choices, conflicting objectives, and incomplete information. Humans can weigh context, interpret political events, and question assumptions that an algorithm would simply encode. Risk management is as much about culture and accountability as it is about prediction.
For now, AI can enhance human judgment, not replace it. The most effective systems combine computational foresight with human skepticism and prudence: machines to detect, humans to decide.

The Human Element of Adaptive Finance

Artificial intelligence is reshaping risk management from a static discipline into a dynamic conversation between data and judgment. By forecasting stress and updating itself as conditions change, AI gives institutions time to act rather than react. The same adaptability that protects banks also safeguards households, pensions, and economies that depend on them.
Still, the future of financial stability will depend not only on smarter algorithms but on wiser governance. Models can illuminate where risk is forming, but only people can decide what to do about it. The union of machine learning and human responsibility marks the beginning of a new phase in finance: one where technology strengthens foresight without replacing accountability.

SOURCES:

- Bank for International Settlements (2025) - Predicting Financial Market Stress with Machine Learning (BIS Working Papers No. 1250) - bis.org/publ/work1250.htm
- Journal of Financial Markets (2024) - AI and financial risk management: A review and agenda - sciencedirect.com/science/article/pii/S1057521924004083
- Monetary Authority of Singapore (2024) - Artificial Intelligence Model Risk Management: Observations from a Thematic Review - mas.gov.sg/.../information-paper-on-ai-risk-management-final.pdf
- University of Strathclyde (2023) - Explainable AI for Financial Risk Management - strath.ac.uk/media/.../Explainable_AI_For_Financial_Risk_Management.pdf
- Semantic Scholar (2023). Ethical and technical challenges of AI in financial supervision - pdfs.semanticscholar.org/124c/.../01a0859c0e471238fe57aec3886e0e46c329.pdf