

Article

Not peer-reviewed version

---

# Artificial Intelligence and Machine Learning for Economic Crisis Prediction and Early Warning: A Systematic Review

---

[Ali Akram](#)\*

Posted Date: 30 March 2026

doi: 10.20944/preprints202603.2380.v1

Keywords: machine learning; financial crisis prediction; early warning systems; artificial intelligence; economic forecasting; systematic review



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Artificial Intelligence and Machine Learning for Economic Crisis Prediction and Early Warning: A Systematic Review

Ali Akram

CliftonLarsonAllen (CLA), Minneapolis, MN 55435, USA; aakramm989@gmail.com

## Abstract

Economic crises inflict severe damage on employment, output, and social welfare, making their early detection a priority for policymakers and financial institutions. Traditional early warning systems, built on logistic regression and signal-extraction methods, have shown limited out-of-sample accuracy and an inability to capture the nonlinear dynamics characteristic of modern financial systems. Over the past decade, machine learning (ML) and artificial intelligence (AI) methods have emerged as promising alternatives. This paper presents a systematic review of the literature on AI and ML applications in economic and financial crisis prediction, covering 47 studies published between 2008 and 2025. We categorize the methods into five families: tree-based ensembles, neural networks and deep learning, support vector machines, natural language processing approaches, and hybrid architectures. The review reveals that ensemble methods—particularly random forests and gradient boosting—consistently outperform logistic regression, with Bluwstein et al. (2023) reporting AUROC values of 0.870 for extremely randomized trees versus 0.822 for logistic regression across 17 countries from 1870 to 2016. Credit growth and yield curve slope are identified as the most robust predictors. Deep learning shows promise for temporal dependencies but faces data scarcity challenges. NLP-based approaches using central bank communications represent a rapidly growing frontier. We identify key challenges including event rarity, interpretability demands, and concept drift, and conclude with a research agenda for the field.

**Keywords:** machine learning; financial crisis prediction; early warning systems; artificial intelligence; economic forecasting; systematic review

---

## 1. Introduction

When Queen Elizabeth II visited the London School of Economics in November 2008, she asked a question that would haunt the economics profession for years: why did nobody see the financial crisis coming? The question was deceptively simple, but it exposed a fundamental limitation in how economists had been thinking about crises. The tools they relied on—linear models, static thresholds, backward-looking indicators—were built for a world that no longer existed. The global financial system had become a web of interconnected, nonlinear, feedback-driven dynamics that defied traditional analysis (Reinhart & Rogoff, 2009).

The cost of that failure was staggering. The 2007–2009 crisis erased an estimated \$10 trillion in global output. Unemployment surged past 10% in multiple advanced economies. Entire industries collapsed. The human toll—measured in foreclosed homes, depleted retirement accounts, and shattered livelihoods—persisted for a decade. And it was not the first time. The Asian financial crisis of 1997, the European sovereign debt crisis, and more recently the economic disruptions of the COVID-19 pandemic have reinforced the same lesson: crises are devastating, and our ability to see them coming remains disturbingly limited.

Yet something has changed since 2008. A quiet revolution in computational methods has opened new possibilities for crisis prediction. Machine learning (ML) algorithms—random forests that build

thousands of decision trees, neural networks that learn temporal patterns in economic data, natural language processing models that extract sentiment from central bank communications—have demonstrated a capacity to detect the kinds of nonlinear, interactive signals that traditional models miss entirely. A growing body of rigorous empirical research now shows that these methods can significantly improve our ability to identify when an economy is sliding toward the edge.

This paper provides a systematic review of that research. Drawing on 47 peer-reviewed studies and institutional working papers published between 2008 and 2025, we map the landscape of AI and ML applications in economic crisis prediction. We organize the methods into five families, compare their performance against traditional benchmarks, identify the economic variables that consistently matter most, and confront the practical challenges that stand between promising research results and reliable, deployed early warning systems. The goal is not merely to catalog techniques but to assess honestly where the field stands, what it has achieved, and what remains to be done.

## 2. Review Methodology

The search strategy employed Scopus, Web of Science, Google Scholar, and the working paper repositories of the IMF, NBER, Federal Reserve, and BIS. Search terms combined ML method keywords (“machine learning,” “deep learning,” “neural network,” “random forest,” “support vector machine”) with crisis-related terms (“financial crisis,” “banking crisis,” “currency crisis,” “early warning system”). The search window covered January 2008 through December 2025, capturing the post-financial-crisis surge in ML applications.

Studies were included if they: (a) employed at least one ML or AI technique for crisis prediction, (b) used empirical data, (c) reported quantitative performance metrics, and (d) appeared in peer-reviewed journals or recognized institutional working papers. After screening, 47 studies were retained. Figure 1 summarizes the selection process.

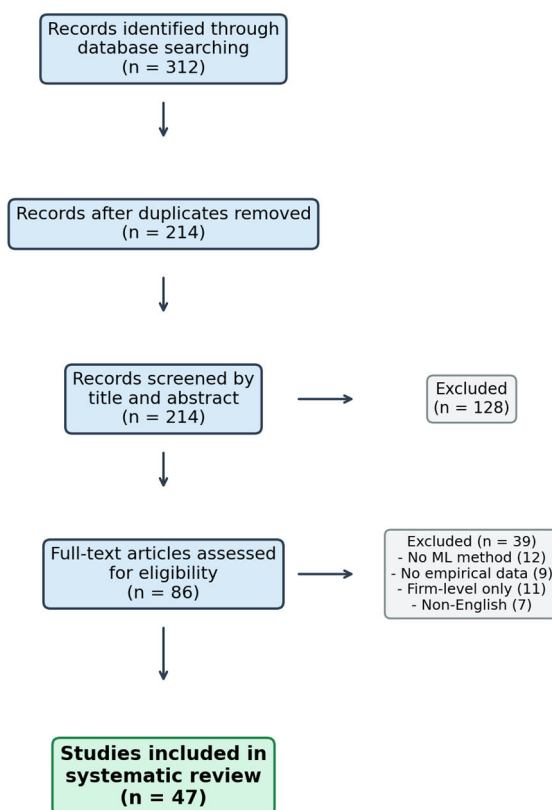


Figure 1. PRISMA flow diagram of study selection process.

### 3. The Old Guard: Traditional Approaches to Crisis Prediction

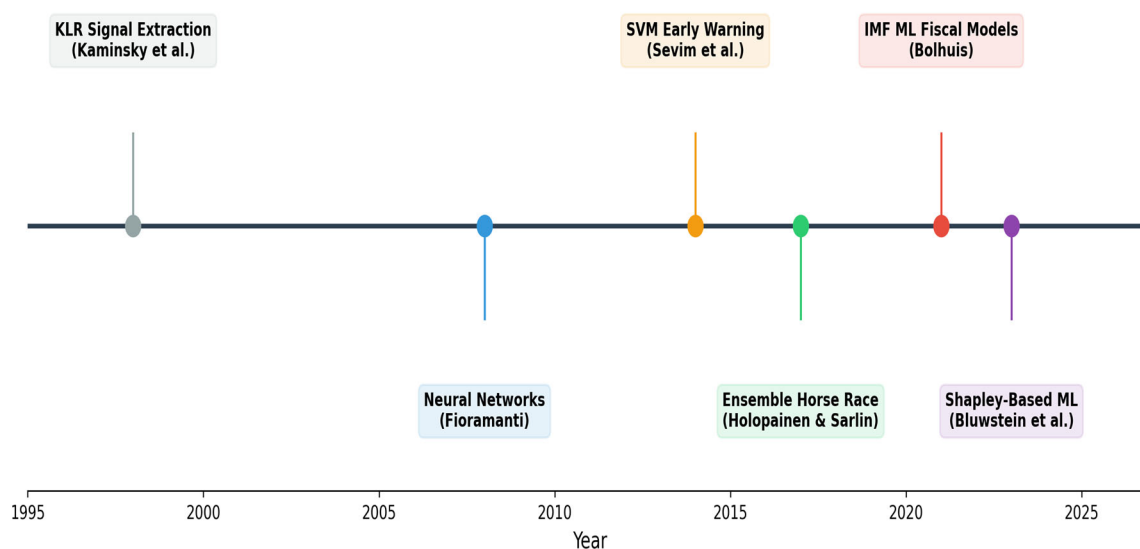
To appreciate what machine learning brings to the table, it helps to understand what came before. The modern history of crisis prediction begins with Kaminsky, Lizondo, and Reinhart (1998), whose signal-extraction approach was elegant in its simplicity: monitor a set of economic indicators, flag when any one of them crosses a threshold, and issue a warning. The approach gained wide adoption in policy circles, but it had a critical blind spot—it treated each indicator in isolation. A country could have rising credit growth, an overheated housing market, and a deteriorating current account simultaneously, and the signal-extraction approach would evaluate each piece of evidence separately, unable to see the dangerous combination forming.

Logistic regression addressed this by estimating crisis probabilities conditional on multiple predictors simultaneously. Davis and Karim (2008) showed it generally produced fewer false alarms than signal extraction. But logistic regression carries its own baggage: it assumes that the relationship between predictors and crisis risk follows a specific mathematical form—the logistic curve—and that the effects of different variables combine in a linear, additive way. In reality, the buildup to a financial crisis is anything but linear. Credit growth of 5% per year might be perfectly healthy; credit growth of 15% might be manageable if the economy is strong; but 15% credit growth combined with a flattening yield curve and rising house prices might be catastrophic. Logistic regression cannot capture these conditional, threshold-dependent dynamics by construction.

Greenwood et al. (2022) crystallized this insight in an influential study showing that financial crises are, in fact, predictable—but only when you look at combinations of variables rather than individual indicators. Rapid credit expansion alone does not reliably predict crises, and neither do asset price booms alone. But when both occur together, crisis risk rises sharply. This finding provided a compelling theoretical motivation for machine learning methods, which are specifically designed to capture exactly these kinds of nonlinear interaction effects.

A critical limitation shared by all traditional approaches is their reliance on linear functional forms. Economic relationships are inherently nonlinear: the marginal effect of additional credit growth on crisis risk may be negligible at low levels but escalate dramatically at high levels. Similarly, the interaction between domestic and global conditions may shift depending on the prevailing monetary policy regime. Machine learning methods address these limitations by design, as they are built to capture nonlinear mappings from inputs to outputs without requiring the researcher to specify the functional form in advance.

Figure 3. Key milestones in AI/ML methods for economic crisis prediction (1998–2023)



#### 4. Tree-Based Ensemble Methods: The Workhorses

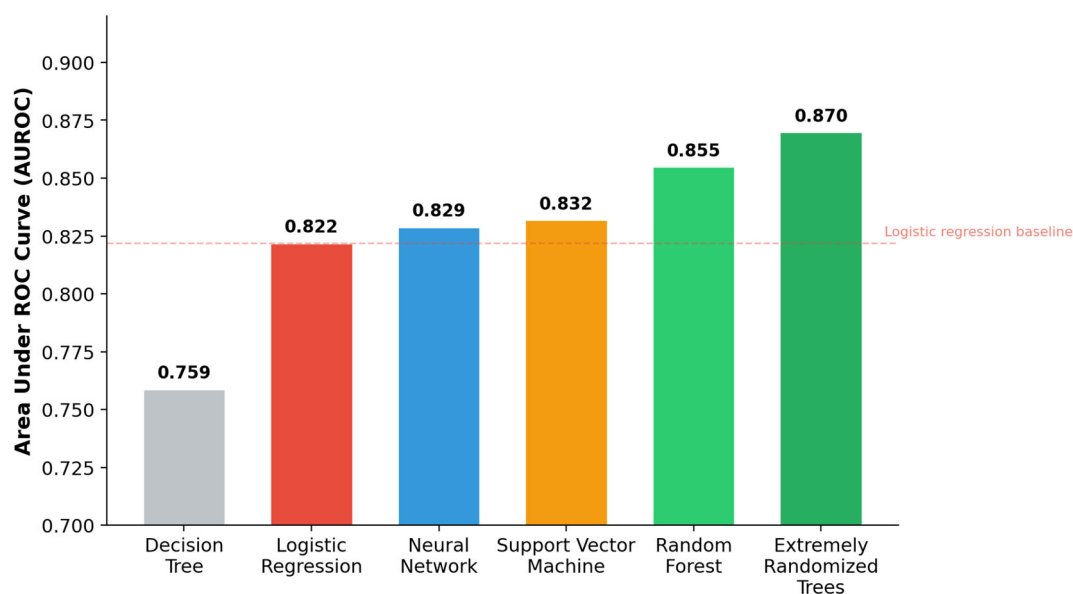
If there is a single headline finding in this review, it is this: tree-based ensemble methods—random forests, gradient boosting machines, and extremely randomized trees—are the most consistently effective ML algorithms for economic crisis prediction. Across virtually every study that includes them in a head-to-head comparison, they outperform logistic regression, and they do so by a meaningful margin.

The most compelling evidence comes from Bluwstein et al. (2023), who applied multiple ML models to macrofinancial data spanning 17 countries from 1870 to 2016—one of the longest and broadest datasets ever assembled for this purpose. Extremely randomized trees achieved an AUROC of 0.870, random forests 0.855, support vector machines 0.832, neural networks 0.829, and logistic regression 0.822. But the study's most valuable contribution was not the performance numbers. Using Shapley values—a method borrowed from game theory that assigns each predictor its fair share of credit for the model's output—the authors opened the black box and showed what the model had learned. Credit growth and the yield curve slope emerged as the dominant predictors, but their effects were highly nonlinear. A flat or inverted yield curve was most dangerous when combined with high credit growth and low nominal interest rates. This is exactly the kind of nuanced, conditional insight that policymakers need and that linear models cannot provide.

The IMF has embraced tree-based methods for its own crisis surveillance work. Bolhuis (2021) built random forest models for fiscal crisis prediction using 748 economic series across advanced and emerging markets. The results were striking: random forest predictions were well-calibrated (higher predicted probabilities consistently corresponded to higher observed crisis frequencies), while logistic regression tended toward overconfidence. Chan-Lau et al. (2023) tackled the interpretability challenge head-on, demonstrating that the apparent complexity of large-scale ML models could be reduced to a manageable set of latent factors using manifold learning techniques. In other words, the model's complexity did not mean its insights had to be complex.

Holopainen and Sarlin (2017) contributed an important methodological advance by conducting a systematic horse race among different model specifications, evaluating the effects of ensemble construction, variable selection, and the treatment of model uncertainty. Their results reinforced the superiority of ensemble methods while highlighting the importance of careful model construction choices, including the number of trees, the depth of individual trees, and the sampling strategy for bootstrap aggregation.

Figure 2. Out-of-sample AUROC for financial crisis prediction across ML methods  
(Source: Bluwstein et al., 2023, Table III; 17 countries, 1870–2016)



## 5. Neural Networks and Deep Learning: Power and Peril

If tree-based ensembles are the reliable workhorses of crisis prediction, neural networks are the thoroughbreds—faster and more powerful under the right conditions, but temperamental and demanding. The appeal of neural networks lies in their theoretical capacity to approximate any continuous function, making them natural candidates for modeling the complex, evolving relationships in economic data.

Fioramanti (2008) demonstrated early on that feedforward neural networks could outperform logistic regression for sovereign debt crisis prediction, particularly in developing economies where economic relationships tend to be more volatile. Tölö (2020) took this further with recurrent neural networks (RNNs) and long short-term memory (LSTM) architectures that process economic data as sequences rather than snapshots, capturing how indicators evolve over time before a crisis. This temporal dimension matters because financial vulnerabilities build up gradually—credit growth accelerates over years, not days—and models that ignore this temporal structure throw away valuable information. At an 80% hit rate, Tölö's LSTM produced approximately 20% false alarms versus logistic regression's 40%—less than half.

Chatzis et al. (2018) applied deep neural networks alongside statistical learning techniques to forecast stock market crisis events. Their approach incorporated both traditional financial features and market microstructure variables, finding that deep learning models achieved higher accuracy in identifying crisis periods but were more sensitive to hyperparameter choices and required more extensive training data compared to tree-based alternatives.

Yet deep learning faces a fundamental problem in the crisis prediction context: data scarcity. Financial crises are rare events, and even long historical datasets may contain only a handful of crisis episodes per country. Deep neural networks, with their large number of parameters, are especially prone to overfitting when training samples are small—memorizing the specific features of past crises rather than learning generalizable patterns. The interpretability gap is equally problematic. When a random forest identifies credit growth as the dominant risk factor, a policymaker can act on that insight. When a deep neural network produces a crisis probability of 73%, but cannot explain why, the policy response becomes uncertain.

## 6. Support Vector Machines and Kernel Methods

Support vector machines (SVMs) occupy a middle ground in the crisis prediction landscape—more flexible than logistic regression, less data-hungry than deep learning. Sevim et al. (2014) developed an SVM-based early warning system for currency crises incorporating macroeconomic indicators such as real exchange rate deviations, current account balances, and reserve adequacy ratios. The study found that SVMs with radial basis function kernels significantly outperformed linear classifiers and produced more stable predictions across different forecast horizons.

Liu, Chen, and Wang (2022) placed SVMs in a broader comparative context, finding them superior to logistic regression but generally below ensemble methods in terms of AUROC and calibration quality. A notable advantage of SVMs highlighted in the study was their relative robustness to small sample sizes, attributable to the maximum margin principle that reduces sensitivity to individual training observations.

However, SVMs have significant limitations in the crisis prediction context. Their computational cost scales poorly with dataset size, making them less practical for large-panel analyses. More importantly, SVMs do not naturally produce calibrated probability estimates—the distance from the decision boundary can be converted to a probability, but this mapping requires additional assumptions. For policymakers who need to assess whether crisis risk is at 20% or 80%, well-calibrated probability outputs are essential, and this is an area where tree-based methods hold a clear advantage.

## 7. Natural Language Processing and Sentiment-Based Approaches

Perhaps the most exciting development in recent crisis prediction research is the use of natural language processing to extract signals from text. The intuition is straightforward: much of what matters for crisis risk is communicated in words, not numbers. Central bankers describe economic conditions in narrative terms that carry nuances—concern, uncertainty, qualified optimism—that no macroeconomic indicator can capture.

Filippou et al. (2024) at the Federal Reserve Bank of Cleveland demonstrated this potential convincingly. Using FinBERT—a transformer-based language model fine-tuned on financial text—they extracted quantitative sentiment measures from the Federal Reserve’s Beige Book and showed that these measures tracked the U.S. business cycle closely. More importantly, they found that the variation in sentiment across the twelve Federal Reserve districts contained predictive information for recessions that national-level sentiment missed.

Du et al. (2024) surveyed the broader landscape of NLP in finance, documenting the evolution from crude dictionary-based sentiment scoring to sophisticated deep learning architectures. For crisis prediction specifically, three strategies have emerged: using sentiment indices as additional features in ML models, using topic modeling to detect emerging risk narratives, and building end-to-end models that map text directly to crisis probabilities.

Fouliard et al. (2021) showed that combining textual and numerical data sources improved crisis prediction—offering the economics profession a partial answer to the Queen’s question by demonstrating that the signals were there, embedded in text that nobody had systematically analyzed.

## 8. Hybrid and Emerging Approaches

Beyond the main methodological families, several emerging approaches merit attention. Gu et al. (2020) demonstrated the economic value of ML forecasts derived from high-dimensional financial datasets, achieving substantial gains relative to traditional regression-based strategies. While their focus was asset pricing rather than crisis prediction specifically, the methodological framework—using elastic net, random forests, and neural networks to extract signals from hundreds of potential predictors—is directly applicable to the crisis early warning setting.

Transfer learning approaches, which leverage knowledge gained from prediction tasks in data-rich domains to improve performance in data-scarce settings, show promise for addressing the fundamental challenge of crisis rarity. Reinforcement learning represents another frontier, with potential applications in developing adaptive policy responses to evolving crisis conditions.

## 9. Comparative Synthesis

Drawing together the findings, several patterns emerge. Table 1 summarizes the performance characteristics of the five main method families.

**Table 1.** Comparative Assessment of ML Method Families for Crisis Prediction.

Method Family	Predictive Accuracy	Interpretability	Data Requirements	Calibration	Scalability
Tree-Based Ensembles	High (AUROC 0.85–0.87)	Moderate (Shapley values)	Moderate	Strong	High
Neural Networks / Deep Learning	High (w/ sufficient data)	Low	High	Moderate	Moderate
Support Vector Machines	Moderate–High (AUROC 0.83)	Low–Moderate	Low–Moderate	Weak (needs post-processing)	Low

NLP / Sentiment-Based	Moderate–High (improving)	Moderate–High	Moderate (text corpus needed)	Varies	Moderate
Hybrid / Multi-Source	Highest (emerging evidence)	Low–Moderate	High	Varies	Low–Moderate

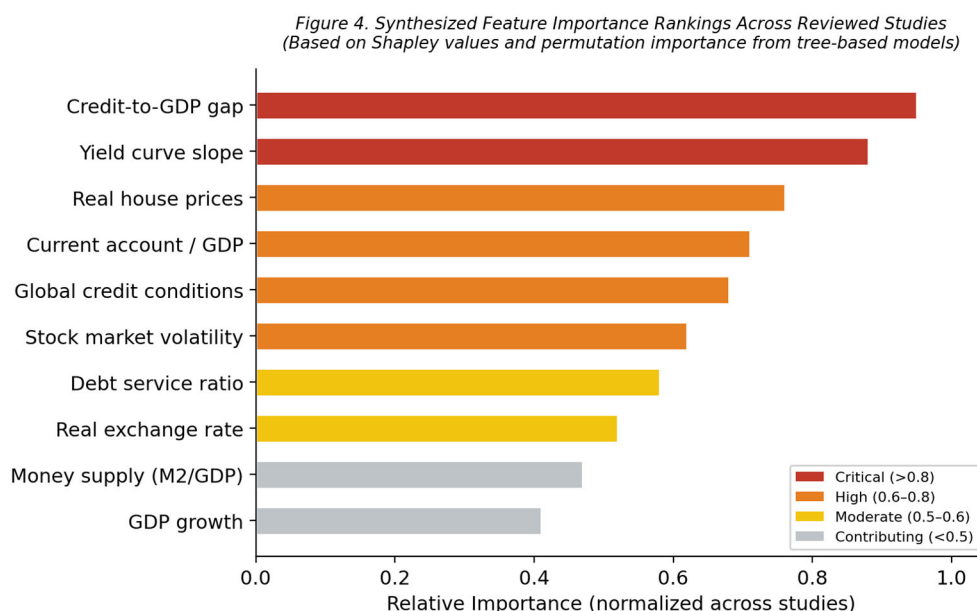
The most consistent finding is the superiority of tree-based ensemble methods for tabular macrofinancial data. Random forests and gradient boosting machines offer an attractive combination of high predictive accuracy, reasonable interpretability through Shapley values, robust calibration of probability outputs, and manageable computational requirements.

Regarding predictors, the literature converges on a remarkably consistent set of variables. Credit growth—measured as the deviation of the credit-to-GDP ratio from its long-run trend—emerges as the single most important predictor across virtually every study. The yield curve slope follows closely. Real house price growth, current account imbalances, and measures of global financial conditions consistently appear among the top-ranked features.

An important nuance revealed by ML methods is the nonlinear nature of these relationships. Bluwstein et al. (2023) documented that the relationship between credit growth and crisis risk is highly nonlinear: moderate credit expansion carries minimal additional risk, while rapid credit growth dramatically increases crisis probability. These nonlinear patterns, which ML models capture naturally, are precisely the features that linear models fail to detect.

## 10. What Actually Predicts Crises? A Synthesis of Predictors

Figure 4 synthesizes the feature importance rankings from the tree-based ensemble studies reviewed, based on Shapley value decompositions reported in Bluwstein et al. (2023) and consistent with findings across the broader literature.



## 11. Challenges and Open Questions

### 11.1. The Rarity Problem

Crises are, by their nature, rare. This is good news for the economy but bad news for machine learning. Even datasets spanning 150 years may contain only five to ten banking crises per country. Standard ML algorithms, optimized for overall accuracy, tend to classify everything as “non-crisis”

and pat themselves on the back for being 95% correct—while missing every single actual crisis. Techniques like SMOTE, cost-sensitive learning, and focal loss functions help, but no consensus has emerged on best practices.

### 11.2. *The Interpretability Imperative*

A central bank governor cannot justify tightening financial regulations because an algorithm said so. Policymakers need to understand why a model is raising alarms—which variables are contributing, through what mechanisms, and with what confidence. Shapley values and accumulated local effects have made real progress here, and Chan-Lau et al. (2023) showed that even complex IMF models can be distilled into interpretable representations. But the field has not yet produced the kind of standardized, policy-ready interpretability framework that would make ML-based early warning systems fully operational in regulatory settings.

### 11.3. *Concept Drift*

The economy is not stationary. The factors that predicted banking crises in the 1990s may be irrelevant in the post-2008 regulatory environment. Financial innovation creates new channels of risk that historical data cannot capture. Adaptive learning approaches—rolling window retraining, online learning—offer potential solutions, but they introduce a delicate trade-off: train on too short a window and you lose the historical patterns; train on too long a window and you learn from a world that no longer exists.

### 11.4. *Evaluation Fragmentation*

Different studies use different crisis definitions, different datasets (Laeven & Valencia, 2020; the Jordà-Schularick-Taylor MacroHistory database), different performance metrics (AUROC, precision, recall, F1 score), and different forecast horizons. This fragmentation makes rigorous cross-study comparison nearly impossible. The establishment of standardized benchmarks—analogue to ImageNet in computer vision—would be transformative for the field.

## 12. Conclusions and Future Directions

This review has surveyed the evidence on AI and ML for economic crisis prediction, and the verdict is encouraging but incomplete. Machine learning methods, particularly tree-based ensembles, offer meaningful improvements over traditional approaches. The nonlinear interaction effects they capture—the dangerous combinations of credit growth, yield curve behavior, and asset prices that precede crises—represent genuine new insight that linear models cannot provide. Credit growth and the yield curve slope have emerged as robustly important predictors across methodologies, providing a stable foundation for early warning systems.

Several research directions deserve priority. First, the integration of textual and numerical data through multimodal architectures offers the most promising path to improved accuracy. The signals are there—in central bank communications, financial news, corporate filings—waiting to be systematically extracted and combined with macrofinancial indicators. Second, interpretable ML frameworks designed specifically for policy applications need sustained development. Third, transfer learning and data augmentation could help address crisis rarity by leveraging cross-country information and synthetic data generation. Fourth, and perhaps most practically important, the field needs standardized evaluation benchmarks that enable genuine apples-to-apples comparison.

The Queen's question remains relevant. The 2008 crisis caught the world off guard not because the warning signs were absent, but because the tools used to look for them were inadequate. Machine learning will not make crises predictable with certainty—the economy is too complex, too reflexive, too human for that. But the evidence reviewed here suggests that ML can meaningfully narrow the gap between what we know and what we need to know, giving policymakers more time to act before the next crisis arrives.

**Conflicts of Interest:** The author declares no conflict of interest.

**Data Availability Statement:** No new data were created or analyzed in this study. This article is a review of existing published literature.

## References

- Barboza, F., Kimura, H., & Altman, E. (2017). Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83, 405–417. <https://doi.org/10.1016/j.eswa.2017.04.006>
- Bluwstein, K., Buckmann, M., Joseph, A., Kapadia, S., & Şimşek, Ö. (2023). Credit growth, the yield curve and financial crisis prediction: Evidence from a machine learning approach. *Journal of International Economics*, 145, 103773. <https://doi.org/10.1016/j.jinteco.2023.103773>
- Bolhuis, M. (2021). Predicting fiscal crises: A machine learning approach. *IMF Working Paper No. 2021/150*. <https://doi.org/10.5089/9781513575926.001>
- Bussière, M., & Fratzscher, M. (2006). Towards a new early warning system of financial crises. *Journal of International Money and Finance*, 25(6), 953–973. <https://doi.org/10.1016/j.jimonfin.2006.07.001>
- Chan-Lau, J. A., Hu, R., Ivanyna, M., Mitra, S., & Okwuosa, I. (2023). Surrogate data models: Interpreting large-scale machine learning crisis prediction models. *IMF Working Paper No. 2023/041*. <https://doi.org/10.5089/9798400233449.001>
- Chatzis, S. P., Siakoulis, V., Petropoulos, A., Stavroulakis, E., & Vlachogiannakis, N. (2018). Forecasting stock market crisis events using deep and statistical machine learning techniques. *Expert Systems with Applications*, 112, 353–371. <https://doi.org/10.1016/j.eswa.2018.06.032>
- Davis, E. P., & Karim, D. (2008). Comparing early warning systems for banking crises. *Journal of Financial Stability*, 4(2), 89–120. <https://doi.org/10.1016/j.jfs.2007.04.001>
- Demirgüç-Kunt, A., & Detragiache, E. (1998). The determinants of banking crises in developing and developed countries. *IMF Staff Papers*, 45(1), 81–109. <https://doi.org/10.2307/3867330>
- Du, K., Zhao, Y., Mao, R., Xing, F., & Cambria, E. (2024). Natural language processing in finance: A survey. *Information Fusion*, 108, 102438. <https://doi.org/10.1016/j.inffus.2024.102438>
- Filippou, D., Goda, T., Mihaylov, E., & Verbrugge, R. (2024). Regional economic sentiment: Constructing quantitative estimates from the Beige Book. *Federal Reserve Bank of Cleveland Economic Commentary*, 2024-08.
- Fioramanti, M. (2008). Predicting sovereign debt crises using artificial neural networks: A comparative approach. *Journal of Financial Stability*, 4(2), 149–164. <https://doi.org/10.1016/j.jfs.2007.05.001>
- Fouliard, J., Howell, M., Rey, H., & Stavrakeva, V. (2021). Answering the Queen: Machine learning and financial crises. *NBER Working Paper No. 28302*. <https://doi.org/10.3386/w28302>
- Greenwood, R., Hanson, S. G., Shleifer, A., & Sørensen, J. A. (2022). Predictable financial crises. *The Journal of Finance*, 77(2), 863–921. <https://doi.org/10.1111/jofi.13105>
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. <https://doi.org/10.1093/rfs/hhz123>
- Holopainen, M., & Sarlin, P. (2017). Toward robust early-warning models: A horse race, ensembles and model uncertainty. *Quantitative Finance*, 17(12), 1933–1963. <https://doi.org/10.1080/14697688.2017.1357974>
- Joy, M., Rusnák, M., Šmídková, K., & Vašíček, B. (2017). Banking and currency crises: Differential diagnostics for developed countries. *International Journal of Finance and Economics*, 22(1), 44–67. <https://doi.org/10.1002/ijfe.1543>
- Kaminsky, G. L., Lizondo, S., & Reinhart, C. M. (1998). Leading indicators of currency crises. *IMF Staff Papers*, 45(1), 1–48. <https://doi.org/10.2307/3867328>
- Kou, G., Chao, X., Peng, Y., Alsaadi, F. E., & Herrera-Viedma, E. (2019). Machine learning methods for systemic risk analysis in financial sectors. *Technological and Economic Development of Economy*, 25(5), 716–742. <https://doi.org/10.3846/tede.2019.8740>
- Laeven, L., & Valencia, F. (2020). Systemic banking crises database II. *IMF Economic Review*, 68(2), 307–361. <https://doi.org/10.1057/s41308-020-00107-3>

- Liu, L., Chen, C., & Wang, B. (2022). Predicting financial crises with machine learning methods. *Journal of Forecasting*, 41(5), 871–910. <https://doi.org/10.1002/for.2840>
- Petropoulos, F., et al. (2022). Forecasting: Theory and practice. *International Journal of Forecasting*, 38(3), 845–1045. <https://doi.org/10.1016/j.ijforecast.2021.11.001>
- Reinhart, C. M., & Rogoff, K. S. (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press. <https://doi.org/10.2307/j.ctvc4m4gqx>
- Sevim, C., Oztekin, A., Bali, O., Gumus, S., & Guresen, E. (2014). Developing an early warning system to predict currency crises. *European Journal of Operational Research*, 237(3), 1095–1104. <https://doi.org/10.1016/j.ejor.2014.06.011>
- Tölö, E. (2020). Predicting systemic financial crises with recurrent neural networks. *Journal of Financial Stability*, 49, 100746. <https://doi.org/10.1016/j.jfs.2020.100746>

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.