Deep Temporal Models for Early Warning of Financial Market Instability

Haley Lara
Department of Computer Science, University of Houston, Houston, TX, USA.
haleylara74@uh.edu

Kuncara Suartini
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
kuncarasuartini@oregonstate.edu

Jared Costa
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
costa1970@uab.edu

Abstract

Financial market instability is increasingly understood as a socio-technical phenomenon whose early warning cannot be reduced to a single indicator, model, or institutional viewpoint. This paper examines deep temporal models as components of early warning systems for financial market instability, with attention to architecture, infrastructure, governance, sustainability, robustness, fairness, and policy. The analysis begins from established early warning research on currency, banking, and sovereign crises and then connects that literature to contemporary temporal representation learning, including recurrent, convolutional, attention-based, and multi-horizon forecasting architectures. The central argument is that predictive performance alone is insufficient for deployment in high-stakes financial environments. Effective systems require calibrated uncertainty, regime awareness, explainable signals, human oversight, and governance mechanisms that clarify accountability across central banks, regulators, market infrastructures, and private institutions. Deep temporal models can detect subtle changes in liquidity, volatility spillovers, correlation structures, and sentiment, but they can also amplify data biases, produce false alarms, and create new systemic dependencies. The paper therefore treats early warning as a decision infrastructure rather than a model artifact. It considers evaluation beyond statistical accuracy, including lead time, false alarm burden, distributional impact, operational resilience, and institutional capacity. Cross-domain comparisons with epidemiology, climate monitoring, and power grid management clarify common trade-offs between sensitivity, specificity, and legitimacy. The conclusion proposes a layered research and policy agenda that integrates temporal modeling with model risk management, regulatory reporting, fairness auditing, and long-term maintenance. Early warning systems should be designed to inform judgment, not replace it, and to preserve financial stability while remaining accountable to democratic and market institutions.

Keywords

financial instability, early warning systems, deep temporal models, systemic risk, regime shifts, AI governance, model risk management, financial regulation.

1. Introduction

Financial market instability rarely emerges as a discrete event with a clear beginning. It typically develops through interacting vulnerabilities, including leverage, liquidity mismatch, concentrated exposures, correlated expectations, and institutional incentives that reward short-term risk-taking. Classic work on twin crises demonstrated that banking and balance-of-payments failures often share common antecedents, making early detection a matter of identifying fragile configurations rather than predicting a single shock [1]. Subsequent early warning research refined this insight by combining macroeconomic indicators, financial market variables, and institutional conditions into probabilistic assessments of crisis likelihood [2]. Later studies of the global financial crisis emphasized that leading indicators must be interpreted in relation to global liquidity conditions, cross-border exposures, and policy credibility [3]. These contributions remain foundational, but they also reveal a persistent limitation: most early warning frameworks were designed for relatively low-frequency data and limited model complexity.

Contemporary financial markets generate data at speeds and granularities that earlier systems could not anticipate. Order book activity, derivatives pricing, repo market conditions, payment flows, news sentiment, and alternative data streams create a dense temporal record of market behavior. At the same time, financial instability is shaped by structural dependencies that are not directly observable in prices, including hidden leverage, collateral chains, and expectations about central bank intervention. Systemic risk research has shown that financial institutions and markets behave less like independent actors and more like an ecosystem in which shocks can propagate through overlapping portfolios and funding relationships [4]. This ecological perspective implies that early warning must capture both temporal dynamics and relational structure. A recent strand of decision intelligence research further argues that trust calibration and multi-agent coordination are necessary when risk signals are distributed across institutions and models with different objectives [5].

Deep temporal models are attractive in this setting because they can learn representations from long sequences, irregular observations, and heterogeneous inputs without requiring all relevant features to be specified in advance. However, the translation from promising model class to deployable early warning infrastructure is difficult. Financial authorities must decide what constitutes an actionable signal, how much false alarm burden is tolerable, how to explain warnings to political and market audiences, and how to assign responsibility when warnings are missed or when they trigger destabilizing reactions. These questions are not peripheral. They determine whether a technically accurate model becomes a legitimate component of macroprudential governance or remains an experimental tool with limited operational influence. This paper therefore examines deep temporal models as parts of a broader system architecture. It emphasizes structural trade-offs, institutional fit, sustainability, fairness, robustness, and policy implications rather than narrow model performance.

2. Conceptual Foundations of Financial Instability and Early Warning

The financial instability hypothesis describes how periods of prosperity can breed fragility as economic units take on more debt and speculative financing becomes normalized [6]. In this view, instability is endogenous: the very conditions that support growth create vulnerabilities that later manifest as distress. A related perspective from complexity economics holds that markets can develop critical states in which small perturbations produce disproportionate outcomes [7]. These conceptual traditions caution against treating instability as an exogenous shock that can be predicted by watching a small set of warning lights. Instead, they suggest

that early warning requires attention to changes in the structure of expectations, financing, and market coordination. Temporal models are well suited to this task because they can detect shifts in the statistical relationships among variables, not only changes in their levels.

Regime-switching approaches provide one bridge between economic theory and temporal modeling. Work on nonstationary time series and business cycles demonstrated that economic relationships can change abruptly across regimes, requiring models that allow parameters or dynamics to differ by state [8]. In financial markets, regimes may correspond to calm periods, volatility clusters, liquidity stress, policy intervention, or crisis contagion. Deep temporal models can learn regime-sensitive representations implicitly, but this flexibility creates interpretive challenges. A model may appear accurate because it has learned to recognize a recurring historical episode rather than because it has identified a causal mechanism. When the next crisis differs in form, the model may fail precisely when it is most needed. This risk is especially pronounced in financial early warning, where the number of major crisis observations is small and the cost of false confidence is high.

Financial instability also has a strong cross-sectional dimension. Volatility spillovers across global equity markets show that shocks do not remain contained within national borders or asset classes [9]. Measures of conditional value at risk and systemic contribution further demonstrate that institutions can be individually stable yet collectively fragile when their distress is correlated with the market as a whole [10]. Deep temporal models can incorporate such dependencies by processing multiple time series jointly, learning attention weights, or combining graph-like relational information with sequential encoders. Yet the resulting signals may be difficult to audit. If a warning arises from complex interactions among hundreds of features, supervisors may struggle to determine whether the signal reflects a genuine vulnerability, a data artifact, or an adaptive market response. This interpretive gap is a governance problem as much as a technical one.

High-risk systems theory adds another caution. Normal accident theory argues that tightly coupled and complex systems are prone to failures that cannot be eliminated entirely through better design [11]. Financial markets exhibit tight coupling through high-speed trading, collateral chains, and real-time settlement, while also displaying complexity through overlapping products and opaque intermediaries. Tail-risk perspectives similarly warn that rare, high-impact events may fall outside the range of historical data and conventional risk models [12]. These insights do not make early warning impossible, but they shift the objective from perfect prediction to resilient anticipation. A useful early warning system should help institutions prepare, communicate uncertainty, and reduce vulnerability even when the precise timing and form of instability cannot be known.

3. Deep Temporal Architectures for Instability Detection

Recurrent neural networks, especially long short-term memory models, were developed to capture long-range dependencies in sequential data and have become a standard baseline for financial time series [13]. Their gate mechanisms allow information to persist across many time steps, which is useful for detecting slowly building imbalances such as rising leverage, deteriorating funding conditions, or persistent divergence between credit and equity markets. In an early warning context, recurrence can model the accumulation of vulnerability rather than treating each observation independently. However, recurrent models can be difficult to train, sensitive to sequence length, and opaque in their internal state. They may also overfit to historical episodes if regularized and validated only on standard predictive metrics.

Attention-based architectures, particularly transformers, changed sequence modeling by allowing models to weigh relationships across time steps without relying solely on recurrent state [14]. In finance, attention can help identify which past events or contemporaneous signals are most relevant to a current warning. For example, a model might learn that funding stress in one market becomes more predictive when accompanied by currency depreciation and deteriorating bank equity prices. Convolutional sequence models offer a complementary approach by capturing local temporal patterns efficiently and often performing competitively with recurrent architectures [15]. Multi-horizon forecasting models such as temporal fusion transformers attempt to combine interpretability, variable selection, and long-range prediction, which aligns with the needs of early warning systems that must produce warnings at different lead times [16].

The choice among these architectures is not merely technical. It affects system latency, computational cost, explainability, and maintenance. A recurrent model may be easier to deploy in a low-latency environment but harder to interpret. A transformer may capture complex dependencies but require substantial data and compute. A convolutional model may be efficient but less suited to very long historical dependencies. Ensemble approaches can improve robustness but increase operational complexity and make accountability more diffuse. In practice, an early warning system may combine multiple temporal models with simpler statistical indicators, expert rules, and scenario analysis. The goal is not to crown a single architecture but to design a portfolio of signals that performs reliably across regimes and can be explained to decision-makers.

Deep temporal models also need to represent uncertainty. Financial instability warnings are inherently probabilistic, and decision-makers need to distinguish between a model that is confident about elevated risk and one that is uncertain because data are sparse or regime conditions are unfamiliar. Bayesian approximations, ensembles, conformal methods, and quantile forecasting can provide uncertainty estimates, but these must be integrated into the operational workflow. A warning without a calibrated confidence level may encourage either complacency or overreaction. A warning with uncertainty but no decision rule may be ignored. The architecture must therefore connect temporal inference to institutional thresholds, escalation procedures, and communication protocols. This connection is where model design meets governance.

4. System Architecture and Data Infrastructure

An early warning system for financial instability is not simply a trained model. It is a data-intensive infrastructure that must ingest, validate, transform, store, and serve heterogeneous information. Market data arrive at different frequencies, from tick-level trades to quarterly regulatory filings. Alternative data may be sparse, proprietary, or unstable in definition. Macroeconomic series are often revised, which creates look-ahead bias if backtests use final revised values rather than the information available at the time. A robust architecture must track data provenance, versioning, and transformation logic so that warnings can be reproduced and audited. Without such infrastructure, sophisticated temporal models may learn from artifacts that would never have been observable in real time.

The system also requires separation between experimentation and production. Research environments can explore new features, architectures, and labels, while production environments must enforce stability, latency, and access controls. A feature store can help maintain consistency between training and inference, but it also becomes a critical dependency. If feature definitions change without coordination, model behavior may drift. If

data pipelines fail silently, warnings may disappear. If third-party data providers change coverage, historical comparisons may become invalid. These are ordinary engineering concerns, but in financial stability contexts they have systemic implications. A central bank or regulator relying on an early warning system must be able to explain not only what the model predicted but also what data it saw and how those data were processed.

Infrastructure choices also affect fairness and inclusion. Models trained primarily on data from large, advanced economies may perform poorly for emerging markets with different market microstructure, reporting standards, and institutional arrangements. If early warning signals influence capital flows, regulatory intervention, or emergency liquidity provision, poor performance in under-represented jurisdictions can create distributive harms. A system designed only for the largest banks may miss risks in smaller institutions that collectively matter during a crisis. Data infrastructure should therefore include coverage assessments, performance disaggregation across jurisdictions and institution types, and processes for incorporating local knowledge. Fairness in financial early warning is not identical to fairness in consumer credit, but it still concerns who bears the costs of false positives, false negatives, and unequal model quality.

5. Governance, Fairness, and Regulatory Implications

The governance of AI in finance must address transparency, accountability, and the distribution of responsibility across public and private actors. Ethical frameworks for AI societies emphasize beneficence, non-maleficence, autonomy, justice, and explicability, but these principles require translation into specific financial processes [17]. Algorithmic ethics research further shows that transparency, bias, and accountability are often entangled: a model may be transparent in code but opaque in its social effects [18]. In early warning, the relevant question is not only whether a model can be explained in technical terms but whether its warnings can be contested, reviewed, and corrected by accountable institutions. Internal auditing frameworks for algorithmic systems provide useful guidance by emphasizing documentation, testing, monitoring, and redress across the lifecycle [19].

Central banks and regulators already operate under mandates that include price stability, financial stability, consumer protection, and market integrity. Deep temporal early warning systems must fit within these mandates rather than existing as separate technical projects. A warning that triggers supervisory action may affect institutions' funding costs, reputations, and market access. A warning that is disclosed too broadly may become self-fulfilling by prompting preemptive withdrawals or asset sales. A warning that is kept secret may undermine legitimacy if later revealed. Governance must therefore specify who sees what signal, under what conditions, and with what opportunity to respond. It must also address conflicts of interest when private institutions contribute data or models to a public early warning system while also being subject to its outputs.

Fairness concerns arise at several levels. First, model errors may be distributed unevenly across countries, sectors, or institution sizes. Second, the burden of compliance may fall disproportionately on smaller institutions that lack the capacity to interpret or contest model outputs. Third, the use of alternative data may advantage institutions with superior data access, thereby reinforcing concentration. Fourth, if early warning signals inform macroprudential tools such as capital buffers or lending restrictions, the incidence of those tools may fall on borrowers and communities rather than on the institutions whose risk-taking created vulnerability. A fair system requires impact assessments, appeal mechanisms, and periodic review of whether the benefits and burdens of early warning are broadly shared.

6. Deployment, Sustainability, and Robustness

Deployment is often treated as the final step after model development, but in financial stability contexts it is a continuous sociotechnical process. Models must be monitored for drift, retrained when relationships change, and retired when they no longer add value. Concept drift is especially challenging because financial regimes can shift for political, technological, or institutional reasons. A model trained before a major regulatory reform may learn relationships that no longer hold. A model trained during a period of unconventional monetary policy may misinterpret the return to conventional policy as instability. Sustainability therefore requires not only computational efficiency but also institutional capacity for maintenance, documentation, and succession planning. If a small team of specialists builds a system and later leaves, the system may become an orphaned dependency.

Robustness must also address adversarial and accidental threats. Market participants may learn to game early warning indicators if those indicators become known. Data poisoning, spoofing, and cyberattacks could distort the inputs on which warnings depend. Even without malice, correlated data failures can create false reassurance or panic. A robust system should include redundancy across data sources, anomaly detection at the pipeline level, scenario stress testing, and manual override protocols. It should also be designed for graceful degradation: if a key data feed fails, the system should communicate reduced confidence rather than produce a misleadingly precise warning. These requirements resemble those in other critical infrastructures, where safety and resilience depend on anticipating failure modes rather than assuming perfect operation [11].

Sustainability also has an environmental and economic dimension. Large temporal models can consume significant energy during training and inference, especially when many institutions develop similar systems independently. Shared infrastructure, federated learning, or public-private data utilities might reduce duplication, but they raise questions about competition, privacy, and control. Financial stability is a public good, yet the data and expertise needed to produce it are often privately held. A sustainable architecture must balance efficiency with pluralism, ensuring that a single dominant model or vendor does not become a new source of systemic vulnerability. The goal is not only to predict instability but to avoid creating fragile dependencies in the prediction infrastructure itself.

7. Case Illustrations and Cross-Domain Comparisons

The 2008 global financial crisis illustrates both the value and limits of early warning. Before the crisis, a range of indicators suggested rising housing leverage, deteriorating underwriting standards, and growing reliance on short-term funding. Yet many official and private models failed to connect these signals into an actionable warning, partly because they underestimated contagion through opaque securitization and derivatives networks [3]. A deep temporal model trained on available data might have detected anomalies in funding markets or correlations across asset classes, but it might also have been blind to off-balance-sheet exposures that were not reported in real time. This case shows that data availability and institutional incentives can matter more than model architecture.

The COVID-19 market turmoil of 2020 offers another instructive example. The shock was exogenous in origin but propagated through financial channels that were already fragile, including Treasury market liquidity, money market funds, and corporate credit. An early warning system focused only on traditional banking indicators would have missed key vulnerabilities. A system incorporating market microstructure, fund flows, and cross-asset

dependencies might have provided useful signals, but it would also have faced a rapidly changing policy environment. Cross-domain comparisons with pandemic early warning highlight similar tensions. Public health surveillance must balance sensitivity and specificity, protect privacy, and communicate uncertainty to diverse populations. Climate early warning systems face long time horizons and deep uncertainty about tipping points. Power grid monitoring must handle high-frequency data and cascading failures. In each domain, the hardest problems are not purely predictive but institutional and communicative.

These comparisons suggest that financial early warning should be designed as a learning system. After each episode of stress or near miss, institutions should review which signals were useful, which were misleading, and how decision-making unfolded. This requires psychological realism about how experts interpret probabilistic warnings. Prospect theory and behavioral decision research show that people evaluate risks relative to reference points and are often asymmetric in their response to gains and losses [25]. A warning that is framed as a probability may be ignored, while one framed as a potential loss may trigger overreaction. The design of dashboards, narratives, and escalation rules should therefore be informed by behavioral research as much as by statistical validation.

8. Evaluation, Uncertainty, and Decision Integration

Evaluating an early warning system requires metrics that reflect operational value. Classification accuracy is insufficient when crises are rare and class imbalance is extreme. A model that predicts calm at all times may achieve high accuracy while being useless. Precision, recall, lead time, false alarm rate, and area under the receiver operating characteristic curve are more informative, but even these metrics do not capture the cost of intervention. A false positive that triggers unnecessary capital restrictions may slow growth and harm borrowers. A false negative that misses a crisis may lead to severe dislocation and public bailouts. Evaluation should therefore include decision curves, cost-sensitive analysis, and scenario-based stress testing that reflect the actual authorities and instruments available.

Uncertainty quantification is central to this evaluation. Temporal models should distinguish aleatory uncertainty, which arises from inherent randomness, from epistemic uncertainty, which arises from limited knowledge or unfamiliar regimes. In practice, this distinction is difficult, but it matters for decision-making. If uncertainty is high because the model has not seen similar conditions, decision-makers may rely more on expert judgment or scenario analysis. If uncertainty is low but risk is elevated, they may consider preemptive action. Ensemble methods, Bayesian neural networks, quantile regression, and conformal prediction can help, but none eliminates the need for human interpretation. The system should present confidence, disagreement among models, and data quality indicators alongside the warning itself.

Decision integration also requires clarity about thresholds and authority. Who decides when a warning becomes an action? Is the action supervisory, monetary, communicative, or procedural? How are disagreements between models and experts resolved? What is the role of private institutions that may receive early warnings? These questions should be answered before a crisis, not during one. A layered protocol can help: automated signals feed into a risk assessment committee, which can escalate to senior decision-makers, who may choose among pre-defined response options. The model informs but does not command. This human-in-the-loop principle is widely endorsed in AI governance, but it must be operationalized with training, resources, and accountability mechanisms [24].

9. Policy and Institutional Design

Policy design for deep temporal early warning should begin with the recognition that financial stability is a public good with cross-border externalities. International organizations have emphasized the need for robust monitoring of systemic risk, liquidity conditions, and cross-border spillovers [20]. Regulatory bodies have also examined the financial stability implications of AI and machine learning, noting both efficiency gains and new vulnerabilities such as herding, procyclicality, and opacity [21]. The International Monetary Fund has repeatedly highlighted the interaction between global shocks, financial conditions, and emerging market vulnerabilities [22]. These institutional perspectives suggest that early warning systems should be developed within existing macroprudential frameworks rather than as isolated innovation projects.

Regulatory technology and supervisory technology can support this integration. FinTech and RegTech research has shown that digital tools can improve reporting, compliance, and risk monitoring, but also raise questions about data standardization, supervisory capacity, and regulatory fragmentation [23]. A public-private data infrastructure might allow regulators to access aggregated, privacy-preserving signals while preserving commercial confidentiality. However, such infrastructure requires clear rules on data use, retention, and access. It also requires investment in supervisory expertise, because regulators cannot govern models they do not understand. Capacity building should include training in temporal modeling, uncertainty communication, and algorithmic auditing, not only for technical staff but also for senior officials who must make decisions under uncertainty.

International coordination is equally important. Financial instability does not respect jurisdictional boundaries, and early warning signals may be most valuable when shared across countries. Yet information sharing raises sovereignty, privacy, and market sensitivity concerns. A tiered approach might distinguish between technical signals, aggregated risk assessments, and actionable warnings. Technical signals could be shared among central banks and international organizations, while public communication is calibrated to avoid panic. Governance should specify how disagreements among national authorities are handled and how warnings are revised. The legitimacy of the system will depend on whether it is seen as technocratic and impartial or as serving particular national or institutional interests.

10. Conclusion

Deep temporal models offer significant potential for early warning of financial market instability. They can learn from high-frequency, heterogeneous, and long-range data; detect subtle shifts in correlation and liquidity; and produce multi-horizon signals that support preemptive action. Yet the central lesson of this paper is that prediction is only one component of an effective early warning system. The harder challenges are architectural, institutional, and normative. Data pipelines must be reliable and auditable. Models must be robust to regime shifts, adversarial manipulation, and concept drift. Uncertainty must be communicated honestly. Governance must assign responsibility, protect fairness, and preserve the capacity for human judgment. Deployment must be sustainable, not dependent on fragile teams or opaque vendors. Policy must integrate early warning into macroprudential frameworks while respecting international coordination and democratic accountability.

The future research agenda should therefore be interdisciplinary. Computer scientists can develop temporal models that are more interpretable, uncertainty-aware, and robust to distributional change. Economists can refine theories of endogenous instability and identify

leading indicators that reflect structural vulnerability rather than historical coincidence. Legal and policy scholars can design governance mechanisms for public-private data sharing and algorithmic accountability. Behavioral scientists can study how warnings are interpreted by officials, market participants, and the public. Practitioners can build the operational infrastructure that turns signals into timely, proportionate, and legitimate action. The objective is not a crystal ball but a resilient warning capability that helps societies prepare for financial instability without creating new forms of fragility.

References

1. Kaminsky, G. L., & Reinhart, C. M. (1999). The twin crises: The causes of banking and balance-of-payments problems. American Economic Review, 89(3), 473–500.
2. Bussiere, M., & Fratzscher, M. (2006). Towards a new early warning system of financial crises. Journal of International Money and Finance, 25(6), 953–973.
3. Frankel, J., & Saravelos, G. (2012). Can leading indicators assess country vulnerability? Evidence from the 2008–09 global financial crisis. Journal of International Economics, 87(2), 216–231.
4. Haldane, A. G., & May, R. M. (2011). Systemic risk in banking ecosystems. Nature, 469(7330), 351–355.
5. Liu, T. (2026). Trust-Calibrated Multi-Agent Decision Intelligence for Financial Risk Management under Market Regime Shifts. Available at SSRN 7307503.
6. Minsky, H. P. (1992). The financial instability hypothesis. The Jerome Levy Economics Institute Working Paper No. 74.
7. Sornette, D. (2003). Why stock markets crash: Critical events in complex financial systems. Princeton University Press.
8. Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica, 57(2), 357–384.
9. Diebold, F. X., & Yılmaz, K. (2009). Measuring financial asset return and volatility spillovers, with application to global equity markets. Economic Journal, 119(534), 158–171.
10. Adrian, T., & Brunnermeier, M. K. (2016). CoVaR. American Economic Review, 106(7), 1705–1741.
11. Perrow, C. (1984). Normal accidents: Living with high-risk technologies. Basic Books.
12. Taleb, N. N. (2007). The black swan: The impact of the highly improbable. Random House.
13. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.
14. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.
15. Bai, S., Kolter, J. Z., & Koltun, V. (2018). An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv. https://arxiv.org/abs/1803.01271

16. Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting, 37(4), 1748–1764.

17. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and Machines, 28(4), 689–707.

18. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2), 1–21.

19. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, 33–44.

20. Bank for International Settlements. (2021). Annual economic report 2021. Bank for International Settlements.

21. Financial Stability Board. (2017). Artificial intelligence and machine learning in financial services: Market developments and financial stability implications. Financial Stability Board.

22. International Monetary Fund. (2022). Global financial stability report: Shockwaves from the war in Ukraine test the financial system's resilience. International Monetary Fund.

23. Arner, D. W., Barberis, J. N., & Buckley, R. P. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. Northwestern Journal of International Law & Business, 37(3), 371–413.

24. Buckley, R. P., Zetzsche, D. A., Arner, D. W., & Tang, B. W. (2021). Regulating artificial intelligence in finance: Putting the human in the loop. Sydney Law Review, 43(1), 43–81.

25. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica, 47(2), 263–291.