

Attention-Based Neural Networks for Cross-Asset Risk Prediction

Bertrand Goncalves

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
goncalvesbertrand@binghamton.edu

Abstract

Cross-asset risk prediction has become a central challenge for financial institutions, regulators, and market infrastructures because modern portfolios are exposed to interconnected shocks that move across equities, fixed income, currencies, commodities, and digital assets. This paper examines attention-based neural networks as a system-level architecture for cross-asset risk prediction. Rather than treating attention merely as a predictive technique, the discussion positions it within broader questions of data infrastructure, model governance, robustness, fairness, sustainability, and policy. Attention mechanisms offer advantages for cross-asset settings because they can dynamically weight heterogeneous signals, learn time-varying dependencies, and integrate market microstructure, macroeconomic, and network-based features. However, their deployment also raises concerns about interpretability, regime sensitivity, computational cost, and the concentration of model risk in a small number of opaque platforms. The paper argues that effective cross-asset risk prediction requires a socio-technical design in which attention models are coupled with governance routines, stress-testing protocols, human oversight, and regulatory reporting standards. It reviews architectural trade-offs, data and feature engineering issues, systemic risk implications, deployment constraints, and evaluation practices. It also considers cross-domain lessons from high-frequency trading, network risk measurement, and AI assurance frameworks. The conclusion is that attention-based systems can strengthen cross-asset risk intelligence only if they are embedded in resilient infrastructures and accountable institutional processes. Predictive accuracy alone is insufficient; legitimacy, contestability, and operational sustainability are equally important for systems that inform capital allocation, margin setting, and prudential supervision.

Keywords

attention mechanisms, cross-asset risk, systemic risk, financial infrastructure, model governance, robustness, fairness, sustainability, AI regulation.

1. Introduction

Cross-asset risk prediction concerns the estimation of losses, dependencies, and contagion pathways across multiple asset classes under uncertain and changing market conditions. Traditional risk models often separate asset classes or assume stable correlation structures, but contemporary markets are characterized by rapid spillovers, liquidity fragmentation, and feedback loops between collateral, leverage, and market making [8], [9]. Attention-based neural networks have emerged as a promising family of models because they can assign dynamic importance to different inputs, time steps, and asset relationships [1], [2]. In cross-asset contexts, this capability is valuable because the relevance of one market for another changes during crises, policy shifts, and liquidity events. A model that learns to attend to the

right signals at the right time may detect vulnerability earlier than static factor models or fixed correlation estimates.

The academic literature on deep learning for finance has grown rapidly, with recurrent architectures, convolutional models, and transformer-style designs applied to forecasting, portfolio construction, and risk measurement [3], [4], [5], [6], [7]. Yet much of this work emphasizes predictive performance rather than institutional fit. For a risk system used by banks, asset managers, clearinghouses, or supervisors, predictive performance is only one dimension. Such systems must also be robust to structural breaks, explainable enough for governance, efficient enough for real-time use, and fair in the sense that they do not systematically disadvantage certain counterparties, markets, or public interests. They must also be sustainable in computational and organizational terms, because continuous retraining and data ingestion can create significant cost and environmental externalities. This paper therefore treats attention-based cross-asset risk prediction as a socio-technical infrastructure problem rather than a purely algorithmic one.

The argument proceeds in several stages. It first defines cross-asset risk as a system-level problem shaped by network effects and regime shifts. It then explains why attention mechanisms are architecturally suited to this domain, while also identifying their limitations. Subsequent sections examine data infrastructure, systemic risk, governance, fairness, deployment, sustainability, robustness, and evaluation. Case illustrations and cross-domain comparisons clarify how attention models might be integrated into financial risk practice. The paper concludes that attention-based neural networks can improve cross-asset risk intelligence, but only when they are embedded within accountable institutions, resilient infrastructures, and adaptive regulatory frameworks [10], [22], [25].

2. Cross-Asset Risk as a System-Level Problem

Cross-asset risk is not simply the sum of individual asset risks. It emerges from correlated exposures, common funding conditions, shared dealer networks, and the rehypothecation of collateral. When one market experiences a shock, losses can propagate through margin calls, forced sales, and liquidity withdrawal [11], [12]. The global financial crisis illustrated how mortgage-related instruments, funding markets, and sovereign balance sheets became tightly coupled. More recent episodes, including the March 2020 market turmoil, showed that even traditionally defensive assets can become sources of liquidity stress when investors seek cash. A cross-asset risk model must therefore represent not only price co-movement but also the mechanisms by which shocks travel.

Systemic risk research provides important conceptual foundations. Network-based measures of connectedness show that financial institutions and markets are linked through exposures that may be invisible in simple correlation matrices [13], [14], [15]. Work on financial networks demonstrates that topology matters: dense interconnections can absorb small shocks but amplify large ones [16]. Attention-based models can, in principle, learn these time-varying topologies from data, but they must be designed with network structure in mind. If attention is treated as a generic weighting scheme without reference to institutional relationships, it may capture spurious co-movement rather than true transmission channels. Thus, cross-asset risk prediction requires a hybrid perspective in which neural attention complements, rather than replaces, economic and network reasoning.

Another system-level feature is the presence of regime shifts. Markets alternate between calm and stressed states, and the statistical relationships that hold in one regime often break down

in another. The 2007 quantitative equity crisis, the Flash Crash, and various liquidity spirals demonstrate how crowded positioning and automated execution can generate sudden, non-linear dynamics [17], [18]. An attention model that learns from historical data may perform well in-sample yet fail when the market enters a new regime. This is not merely a technical limitation; it is a governance challenge. Institutions must decide how much confidence to place in a model whose attention weights may reflect past conditions rather than future vulnerabilities. Consequently, cross-asset risk prediction must be evaluated through stress scenarios, reverse stress tests, and expert judgment, not only through backtesting.

3. Attention Mechanisms and Neural Architectural Foundations

Attention mechanisms were originally developed to improve sequence-to-sequence learning by allowing models to focus on relevant parts of an input [2]. The transformer architecture replaced recurrence with self-attention, enabling parallel processing and long-range dependency modeling [1]. In financial risk prediction, this architecture is attractive because asset returns, order book features, macroeconomic releases, and news signals arrive at different frequencies and with different levels of reliability. Attention can learn to prioritize informative signals while down-weighting noise. It can also model cross-sectional relationships by attending across assets rather than only across time. This makes attention a natural candidate for cross-asset risk, where the central question is which markets, institutions, or indicators matter for a given exposure.

However, attention architectures introduce structural trade-offs. Deep transformer models require substantial data and computation, and they can overfit to historical patterns if regularization and validation are weak. Financial time series are low signal-to-noise environments, so a model with millions of parameters may memorize transient correlations [7]. Recurrent models and simpler architectures may sometimes generalize better in small-data regimes [3], [4]. Attention also produces weights that are often interpreted as importance scores, but these weights are not necessarily causal explanations. An attention weight may reflect a model's internal optimization rather than an economically meaningful transmission channel. For risk governance, this distinction is critical: a model that appears interpretable may still mislead decision-makers if its attention map is treated as an explanation without further validation.

Architectural choices also affect deployment. A cross-asset risk engine may need to update intraday, end-of-day, or on-demand during stress events. Full self-attention across hundreds of assets and long histories can be computationally expensive, creating latency and cost challenges for real-time risk management. Hierarchical attention, sparse attention, and factorized designs can reduce cost, but they may lose some cross-asset richness. Ensemble approaches that combine attention networks with gradient-boosted trees, factor models, or network measures can improve robustness, though they complicate governance and attribution. The appropriate architecture is therefore not a universal optimum but a contingent choice shaped by data availability, latency requirements, regulatory expectations, and the institution's risk appetite.

4. Data, Features, and Market Microstructure Integration

Attention-based cross-asset risk models depend on data infrastructure that is broader than traditional risk data warehouses. They may ingest prices, returns, volatilities, credit spreads, yield curves, option-implied measures, order book statistics, trade-level data, fund flows, news sentiment, and macroeconomic indicators [4], [5], [6]. Each data source has different

frequencies, quality characteristics, and legal restrictions. Market microstructure data can reveal liquidity conditions and stressed order flow, but it is voluminous and noisy [18]. Macroeconomic data can provide context, but it is released with lags and revisions. News and alternative data can capture emerging risks, yet they raise issues of bias, licensing, and privacy. A robust system must therefore include data lineage, quality monitoring, and time-stamp governance.

Feature engineering remains essential even with attention models. Raw prices alone may not capture leverage, funding costs, or collateral quality. Cross-asset features such as roll yield, basis spreads, implied volatility skew, and correlation surprises can improve predictive power [11], [12]. Network features that measure connectedness, centrality, and clustering can help the model attend to systemic channels rather than mere price co-movement [13], [14], [15]. Regime indicators, including volatility states, liquidity proxies, and policy uncertainty measures, can condition attention on the prevailing environment. The challenge is to integrate these features without producing a model that is too complex to validate. Feature governance should document why each input is included, how it is transformed, and what risks arise if it becomes unavailable or manipulated.

Data infrastructure also has geopolitical and operational dimensions. Cross-asset risk prediction often requires global market data, which may be subject to localization laws, sanctions, and vendor concentration. If a small number of cloud providers or data vendors control critical inputs, the financial system becomes exposed to correlated operational failures [24]. Attention models can exacerbate this concentration because they require large-scale storage and computation. Institutions should therefore consider data redundancy, on-premises fallback, open standards, and contractual protections. The goal is not to avoid concentration entirely, which may be impractical, but to design a system that remains functional under degraded conditions. In this sense, data infrastructure is not a back-office concern; it is a strategic component of cross-asset risk resilience.

5. Cross-Asset Dependencies, Regime Shifts, and Systemic Risk

Attention models can learn cross-asset dependencies that change over time. During normal markets, attention may focus on carry, momentum, and macro factors. During stress, it may shift toward liquidity, funding, and contagion signals. This dynamic weighting is attractive for systemic risk, because the relevant transmission channels are state-dependent. Network measures of connectedness show that spillovers intensify during crises and that the direction of influence can reverse quickly [13], [14], [15], [16]. Attention can complement these measures by identifying non-linear and higher-order interactions that traditional econometric models may miss. Recent work on trust-calibrated multi-agent decision intelligence for financial risk management under market regime shifts illustrates the growing interest in architectures that adapt to changing regimes while maintaining reliability [10]. This line of inquiry is relevant because cross-asset risk prediction is not only a forecasting problem but also a coordination problem among models, humans, and institutions.

Nevertheless, attention models face serious regime-shift risks. If a model is trained predominantly on low-volatility periods, it may underweight tail dependencies. If it learns from a crisis, it may overfit to the specific policy responses of that episode. Attention weights can be unstable across retraining windows, making it difficult to distinguish genuine changes in market structure from estimation noise. This instability can create procyclical behavior if many institutions use similar models and react to similar signals. The 2007 quantitative equity crisis and the Flash Crash demonstrate how correlated algorithmic responses can amplify

stress [17], [18]. A cross-asset risk system must therefore include circuit breakers, human override protocols, and diversity of models to avoid herding.

Systemic risk also raises questions of boundary and jurisdiction. Cross-asset exposures often cross legal entities, currencies, and regulatory regimes. An attention model trained on global data may identify a vulnerability that is not actionable under any single regulator's mandate. Conversely, a model optimized for one jurisdiction may miss offshore or shadow banking channels. Effective governance requires information sharing among supervisors, central banks, and market infrastructures, but such sharing must respect confidentiality and competition law. Attention models can support supervisory analysis by summarizing complex dependencies, but they cannot resolve the underlying institutional fragmentation. The policy challenge is to build a common analytical language without centralizing sensitive data in a way that creates new risks.

6. Model Governance, Fairness, and Regulatory Implications

Governance is central to attention-based cross-asset risk prediction because these models can influence capital, margin, lending, and investment decisions. A model that is inaccurate may cause losses; a model that is opaque may escape accountability; a model that is biased may allocate risk in ways that harm certain groups or markets. Fairness in finance is often discussed in credit scoring, but it also applies to risk models that affect access to liquidity, funding costs, and market participation. If attention weights systematically disadvantage small banks, emerging markets, or non-bank intermediaries, the model may reinforce existing inequalities [19]. Fairness cannot be reduced to a single metric; it requires deliberation about distributive consequences and the legitimacy of the decision process.

Regulatory frameworks for AI and financial risk are evolving. Principles-based guidance from international bodies emphasizes accountability, transparency, human oversight, and robustness [22], [23], [25]. Model documentation practices such as model cards and datasheets for datasets can improve traceability and contestability [20], [21]. For attention models, documentation should include the training data, feature construction, attention stability tests, known failure modes, and the scope of permissible use. Regulators may require explainability not only of individual predictions but also of the model's overall behavior across scenarios. However, explainability should not become a checkbox exercise. A model can be documented thoroughly and still be unsafe if its assumptions are wrong or if it is deployed outside its validated domain.

Fairness and governance also involve power asymmetries. Large institutions may have the data, talent, and compute to build sophisticated attention models, while smaller institutions rely on third-party vendors. This concentration can create systemic dependence on a few model providers [24]. Regulators may need to supervise critical third parties, require model portability, or mandate common standards for risk reporting. At the same time, overly prescriptive rules can freeze innovation and reduce the diversity of approaches that contributes to system resilience. The governance challenge is to balance standardization with experimentation. Attention-based cross-asset risk prediction should be developed in a way that is auditable, contestable, and aligned with public interest, not merely profitable for early adopters.

7. Deployment, Infrastructure, and Sustainability

Deployment of attention-based risk systems requires more than a trained model. It requires data pipelines, feature stores, model registries, monitoring dashboards, incident response

procedures, and integration with trading, treasury, and risk workflows. Latency requirements vary by use case. Intraday margin calculation may demand near-real-time inference, while strategic risk assessment may tolerate overnight batch processing. Cloud deployment offers elasticity but can introduce vendor lock-in, data residency issues, and correlated outages [24]. On-premises deployment offers control but may lack scalability. Hybrid architectures can balance these trade-offs, but they increase operational complexity. Institutions must design for graceful degradation, so that risk reporting continues even if a model or data feed fails.

Sustainability is an often-overlooked dimension. Large attention models consume significant energy during training and inference, and continuous retraining can multiply that cost. Financial institutions face pressure to reduce carbon footprints, yet risk systems are often exempted from efficiency scrutiny because they are considered mission-critical. This exemption is understandable but incomplete. A cross-asset risk model that is computationally wasteful may be economically and environmentally unsustainable, especially if it is retrained frequently without clear benefit. Techniques such as model distillation, pruning, sparse attention, and scheduled retraining can reduce cost. More importantly, institutions should question whether a complex model is necessary for every task. In some cases, a simpler benchmark may provide sufficient risk insight with lower operational and environmental burden.

Sustainability also has an organizational meaning. A model that requires scarce data scientists to maintain may be fragile if those experts leave or if budgets contract during a downturn. Documentation, code review, and modular design can improve maintainability. Vendor dependence can reduce internal capability, making institutions unable to diagnose failures. A sustainable cross-asset risk system is one that can be understood, maintained, and adapted over time by a resilient team. This is a governance and workforce issue as much as a technical one. Attention models should be introduced with a realistic assessment of long-term operational capacity, not only short-term predictive gains.

8. Robustness, Stress Testing, and Evaluation

Robustness is the ability of a model to maintain acceptable performance under stress, adversarial conditions, and structural change. Attention-based cross-asset risk models can be fragile in several ways. They may be sensitive to input scaling, missing data, or outliers. They may latch onto spurious correlations that disappear out-of-sample. They may be vulnerable to adversarial manipulation if market participants learn how the model reacts to certain signals. Evaluation must therefore go beyond standard backtesting. Stress testing should include historical crises, hypothetical scenarios, reverse stress tests, and sensitivity analysis on attention weights [7], [8], [9]. If small changes in inputs produce large changes in risk estimates, the model is not robust enough for critical use.

Evaluation should also consider stability over time. A model that performs well on average but produces unstable attention patterns may be unreliable for decision-making. Institutions can track attention entropy, rank stability of important features, and turnover in predicted risk rankings. They can compare attention-based predictions with network measures, factor models, and expert assessments [14], [16]. Divergence between models is not necessarily a failure; it can be a signal that the market is in a transitional regime. However, divergence must be managed through governance rather than ignored. Ensembles and model risk overlays can reduce reliance on any single architecture.

Robustness also depends on the broader market ecology. If many firms use similar attention models, their collective behavior may create new vulnerabilities. The 2007 quant crisis and the Flash Crash illustrate how correlated strategies can amplify shocks [17], [18]. Regulators may need to monitor model diversity, leverage, and liquidity provision during stress. Attention models can help identify crowded trades and hidden dependencies, but they can also contribute to crowding if their signals become widely adopted. A resilient system requires not only better models but also limits on leverage, transparent margining, and circuit breakers that interrupt feedback loops.

9. Case Illustrations and Cross-Domain Comparisons

Several episodes illustrate the potential and limits of cross-asset risk prediction. The global financial crisis showed how mortgage credit, funding markets, and sovereign risk became interlinked through collateral and leverage [11], [15]. A cross-asset attention model might have detected rising connectedness between housing-related instruments and short-term funding, but it would also have faced regime-shift challenges because the crisis involved novel institutional failures. The March 2020 turmoil demonstrated how a health shock could trigger simultaneous stress across equities, bonds, commodities, and currencies [12]. Attention models that incorporate liquidity and volatility features could help explain why diversification failed, but they would need real-time data and robust governance to be useful under extreme uncertainty.

The Flash Crash and the 2007 quantitative equity crisis provide microstructural lessons [17], [18]. Automated execution, high-frequency trading, and correlated strategies can transmit shocks rapidly across assets. Attention models that incorporate order book and trade-level features may identify liquidity imbalances earlier, but they must be carefully validated because market microstructure is highly non-stationary. Cross-domain comparisons with other critical infrastructures are instructive. Power grids, air traffic control, and telecommunications networks also rely on automated systems that must fail safely. They use redundancy, human oversight, incident reporting, and stress testing. Financial risk systems could adopt similar principles, especially when attention models are used for real-time decisions [5], [6].

Another comparison is with medical AI and autonomous systems, where fairness, explainability, and post-market surveillance are active policy issues [19], [20]. Financial risk models differ because their predictions can directly affect markets through trading and margining, but the governance lessons are similar. Documentation, audit trails, and monitoring after deployment are essential. Cross-asset risk prediction also has global public goods characteristics: systemic stability benefits everyone, but individual firms may underinvest in resilience. International coordination and standards can help prevent a race to the bottom in model governance [22], [25].

10. Limitations and Future Research

Attention-based cross-asset risk prediction is not a solution to all risk management problems. It cannot eliminate uncertainty, predict unprecedented events with certainty, or resolve political conflicts over regulation. It may inherit biases from historical data, especially if past crises excluded certain market participants or instruments. It may be computationally intensive and difficult to explain. It may also create new forms of model risk if institutions over-rely on attention weights as causal explanations. Future research should develop methods for causal attention, regime-aware validation, and uncertainty quantification in

financial time series. It should also study the systemic effects of widespread adoption, including herding, liquidity, and market stability [10], [24].

Interdisciplinary work is needed. Computer scientists can improve architecture and efficiency, economists can model incentives and contagion, legal scholars can clarify liability and disclosure, and policymakers can design oversight mechanisms. Fairness research should examine how risk models affect access to finance and market liquidity [19]. Sustainability research should measure the energy and organizational costs of large financial AI systems [21], [22]. Regulatory research should explore how to supervise third-party model providers and cross-border data flows [23], [25]. The goal is not to slow innovation but to ensure that innovation is directed toward resilient and legitimate risk intelligence.

Future systems may combine attention networks with graph neural networks, reinforcement learning, and human-in-the-loop decision support. They may use federated learning to share risk insights without centralizing sensitive data. They may incorporate real-time supervisory feeds and market-wide stress indicators. However, technical sophistication must be matched by institutional capacity. A model is only as good as the decisions it informs and the accountability structures that surround it. Cross-asset risk prediction should therefore be evaluated as part of a broader socio-technical system, not as an isolated algorithm.

11. Conclusion

Attention-based neural networks offer a powerful approach to cross-asset risk prediction because they can learn dynamic dependencies, integrate heterogeneous data, and adapt to changing market conditions [1], [8]. They can complement traditional risk models and network measures, providing richer representations of contagion and liquidity risk. However, their benefits are contingent on architecture, data quality, governance, and deployment context. Poorly designed attention models can overfit, obscure causal mechanisms, and create new systemic vulnerabilities. The same flexibility that makes them attractive also makes them difficult to validate and control. For these reasons, attention-based cross-asset risk prediction should be treated as a system-level capability, not a standalone predictive tool.

The paper has argued that robustness, fairness, sustainability, and policy legitimacy are as important as predictive accuracy. Institutions must build data infrastructures that are reliable and contestable, governance routines that document and challenge model behavior, and deployment strategies that allow graceful degradation. Regulators must develop standards for model documentation, stress testing, third-party oversight, and cross-border coordination [22], [23], [25]. Researchers must study regime shifts, causal attention, and the systemic effects of widespread model adoption, including the risks highlighted by recent work on trust-calibrated multi-agent decision intelligence [10]. Ultimately, the value of attention-based cross-asset risk prediction will be measured not only by its ability to forecast losses but by its contribution to a more stable, accountable, and equitable financial system.

References

1. 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.
2. Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. *arXiv*. <https://arxiv.org/abs/1409.0473>

3. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
4. Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.
5. Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial markets: Perspectives from deep learning. *Quantitative Finance*, 19(9), 1449–1459.
6. Zhang, Z., Zohren, S., & Roberts, S. (2020). Deep learning for portfolio optimization. *The Journal of Financial Data Science*, 2(4), 8–20.
7. López de Prado, M. (2018). *Advances in financial machine learning*. Wiley.
8. Hull, J. C. (2018). *Risk management and financial institutions* (5th ed.). Wiley.
9. McNeil, A. J., Frey, R., & Embrechts, P. (2015). *Quantitative risk management: Concepts, techniques and tools* (Revised ed.). Princeton University Press.
10. Liu, T. (2026). Trust-Calibrated Multi-Agent Decision Intelligence for Financial Risk Management under Market Regime Shifts. Available at SSRN 7307503.
11. Bouchaud, J.-P., Farmer, J. D., & Lillo, F. (2009). How markets slowly digest changes in supply and demand. In T. Hens & K. R. Schenk-Hoppé (Eds.), *Handbook of financial markets: Dynamics and evolution* (pp. 57–160). North-Holland.
12. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2), 223–236.
13. Diebold, F. X., & Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1), 119–134.
14. Barigozzi, M., & Brownlees, C. (2019). NETS: Network estimation for time series. *Journal of Applied Econometrics*, 34(3), 347–364.
15. Billio, M., Getmansky, M., Lo, A. W., & Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3), 535–559.
16. Acemoglu, D., Ozdaglar, A., & Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2), 564–608.
17. Khandani, A. E., & Lo, A. W. (2011). What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets*, 14(1), 1–46.
18. Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The Flash Crash: High-frequency trading in an electronic market. *Journal of Finance*, 72(3), 967–998.
19. Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning: Limitations and opportunities*. MIT Press.
20. Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229.

21. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92.
22. National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). NIST AI 100-1.
<https://doi.org/10.6028/NIST.AI.100-1>
23. Financial Stability Board. (2017). Artificial intelligence and machine learning in financial services: Market developments and financial stability implications. Financial Stability Board.
24. Bank for International Settlements. (2021). Regulating big techs in finance. BIS Bulletin No. 45. Bank for International Settlements.
25. OECD. (2021). Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449. OECD Publishing.