

Article

Not peer-reviewed version

Model Monoculture Risk: Systemic AI Convergence in Banking and Financial Markets

[Victor Frimpong](#)*

Posted Date: 5 March 2026

doi: 10.20944/preprints202603.0393.v1

Keywords: artificial intelligence; systemic risk; model monoculture risk; AI convergence; financial stability; banking and financial markets



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

Model Monoculture Risk: Systemic AI Convergence in Banking and Financial Markets

Victor Frimpong

SBS Swiss Business School, Zurich, Switzerland; v.frimpong@research.sbs.edu

Abstract

Artificial intelligence (AI) is significantly changing the landscape of banking and financial markets, improving predictive precision, operational productivity, and decision-making speed. However, underlying these improvements is an underexplored structural vulnerability: the risk of systemic AI convergence. As financial institutions become more dependent on similar foundational models, cloud infrastructure, data suppliers, and AI middleware platforms, diverse firms may begin to interpret and respond to market signals in similar ways. This paper presents the notion of model monoculture risk—a systemic fragility that arises when model similarities, market synchronisation, and infrastructure concentration align. Rather than focusing solely on traditional firm-level model risk management, this paper introduces the M³ Framework (Model–Market–Middleware) to clarify how shared AI architectures can transform localised optimisations into widespread amplification. The Model layer encapsulates the convergence of foundational architectures and training interdependencies; the Market layer reflects synchronised adjustments in portfolios, credit, and liquidity; and the Middleware layer emphasises the concentration of cloud services and vendors that quickens the spread. Collectively, these layers create multiplied exposure to correlated failures. To put this concept into practice, the paper proposes a qualitative Model Monoculture Risk Index (MMRI) to evaluate cross-layer alignment and the erosion of diversity within AI-driven financial systems. By redefining AI governance as a challenge of structural diversification rather than just a validation task, this contribution highlights cognitive diversity as an essential element of financial stability in the era of AI.

Keywords: artificial intelligence; systemic risk; model monoculture risk; AI convergence; financial stability; banking and financial markets

JEL Classification: G21; G23; O33; G28; G01

1. Introduction

Artificial intelligence (AI) is now a central component of finance, integrated into banking, capital markets, and investment management. Initially, AI adoption focused on firm-level automation—boosting efficiency, reducing costs, speeding up compliance, and improving customer service with chatbots and personalised analytics. Today, however, AI's disruptive impact is evolving from isolated applications to more interconnected systems that link behaviours across multiple institutions and markets. This paper highlights that the next phase of AI-driven disruption in finance involves system-level coupling, where shared AI tools, vendors, data pipelines, and model families streamline decision-making across independent institutions. This can lead to correlated errors and heighten financial stability risks (FSB, 2025; Crisanto et al, 2024; Kholia, 2025).

1.1. AI as a Productivity Engine and a Structural Disruptor

AI is transforming finance by enhancing fraud detection, speeding up credit decisions, improving risk monitoring, and scaling advisory services (Aldasoro et al., 2025). Regulatory bodies recognise that AI boosts efficiency and personalisation in products while changing client interactions

and risk management (Perez-Cruz et al., 2025; Crisanto et al, 2024). However, increased efficiency can lead to quicker market adjustments, making them more vulnerable. As predictive models become standardised, unique advantages diminish, leading to crowded market positions. Additionally, embedding AI into workflows (such as credit assessments, surveillance, and client onboarding) shifts decision-making more toward machine and external system inputs, reducing individual organisations' influence.

AI is transforming the financial sector beyond just automation. It changes how competition is structured, redefines human roles in finance, and alters how quickly institutions respond in markets. Recent reports estimate significant workforce disruption in European banking due to AI and digitisation, affecting not only front-office functions but also core operational areas (Foy, 2025).

1.2. *The Overlooked Shift: From Idiosyncratic Adoption to Shared Infrastructures*

This transition highlights the shift from firm-level automation to system-level integration. Previously, AI decisions were based on organisation-specific data, models, and governance, so errors and biases were mostly contained within a single organisation. Now, institutions are increasingly using interconnected systems, such as shared cloud services, third-party analytics, data marketplaces, model development platforms, and common foundation models that are fine-tuned or enhanced through retrieval.

Financial regulators and standard setters are increasingly recognising that vulnerabilities in financial systems are not just due to internal models, but also stem from third-party dependencies and market correlations as AI adoption grows (FSB, 2025). The Financial Stability Board (FSB) highlights the need to monitor AI adoption and its associated vulnerabilities, specifically noting third-party dependencies and market correlations as potential risks to financial stability (FSB, 2025). Similarly, the International Organisation of Securities Commissions (IOSCO) emphasises that AI is being integrated into capital markets in ways that directly affect investor outcomes and market functionality, with correlated behaviour and feedback loops becoming significant (Kholia, 2025).

System-level coupling can occur via various mechanisms:

- **Technology coupling:** Many companies use models based on the same foundation models or vendor components, which can lead to shared blind spots, biases, and failure modes, even after fine-tuning.
- **Data coupling:** Institutions use similar alternative data sources, vendor-curated datasets, or standardised feature stores to align signals and outcomes.
- **Operational coupling:** Workflow tools for surveillance, KYC triage, research summarisation, customer communications, and credit pre-screening standardise decision-making and escalation processes.
- **Governance coupling:** Compliance-driven best practices can lead to reliance on the same control templates and vendor solutions, which limits diversity in risk management strategies.

The result is that the financial system shifts from being just a group of independent decision-makers to a network of institutions that work together cohesively.

1.3. *Why Coupling Matters: Correlated Failure, Crowding, and "Model Monoculture Risk"*

In finance, systemic fragility often arises from correlated exposures and shared responses rather than the failure of a single institution. AI can worsen this issue by standardising what is seen as "rational" behaviour in response to similar signals. When multiple institutions receive the same AI-driven assessments on credit risk, counterparty stress, liquidity, or asset sentiment, they may all react by rebalancing, hedging, deleveraging, or withdrawing liquidity simultaneously.

The paper introduces the term "**model monoculture risk**," referring to *the systemic vulnerability that arises when a sector relies too heavily on a limited number of models, vendors, or AI systems*. This concept aligns with regulators' concerns about the impacts of AI adoption on stability and the challenges they face due to uncertainty and lack of visibility (BIS, 2025). Notably, monoculture risk can arise not only

from identical models but also from shared dependencies, such as common cloud services, middleware, APIs, model families, and benchmarking practices that influence model choices.

The coupling problem has become a geopolitical and strategic issue. Europe's reliance on external AI supply chains is seen as a vulnerability, prompting calls for diversification and a minimum level of domestic capability. In the financial services sector, this translates into a need for operational resilience and for addressing concentration risk.

1.4. Regulation as Both Stabiliser and Accelerator of Convergence

A further challenge arises from the interplay between the spread of AI technology and regulatory compliance requirements. For example, the EU AI Act classifies certain financial applications (particularly those related to creditworthiness and credit scoring) as "high-risk," thereby raising expectations for documentation, governance, and oversight (EBA, 2020). While this may enhance baseline standards and mitigate careless use, it can also lead to a convergence effect: under pressure to comply, organisations may opt for pre-certified vendor solutions and uniform tools, thereby increasing their dependence on a limited range of providers and architectures. In other words, regulatory compliance can reduce diversity in governance but may also constrain technology choices, potentially amplifying the risks associated with monocultures.

At the same time, discussions among regulators are increasingly focused on scaling explainability and accountability (Boeddu et al., 2025). Work by the BIS/FSI indicates that insufficient explainability remains a challenge for regulators and supervisors, particularly as AI is integrated into customer interactions, risk management, and operational decision-making (Perez-Cruz et al., 2025; Crisanto et al., 2024). However, having explainability alone does not address the issue of coupling: a well-explained synchronised outcome can still pose systemic risks if it leads to uniform actions throughout the market.

1.5. Paper Contribution

This introduction presents AI in finance as a structural disruptor, highlighting that its primary emerging risk is system-level coupling, not isolated model failure. It contributes by:

1. Characterising model monoculture risk as a specific vulnerability at the system level influenced by common infrastructures and concentration of vendors.
2. Changing the focus from risk management of the model at the firm level to the ways in which AI implementation can create market interconnections and reliance on third parties.
3. Laying the foundation for a practical measurement method (described later in the paper) that risk leaders and supervisors can utilise to identify early indications of monoculture and coupling.

2. What the Literature Misses: From Model Risk to Model Monoculture

The current academic and supervisory approach to "AI risk in finance" is primarily focused on models and firms. The main concern is whether a specific institution's model is accurate, explainable, well-governed, and validated throughout its lifecycle. This focus is evident in model risk management (MRM) guidelines, which emphasise rigorous development standards, independent validation, outcome testing, and continuous monitoring. U.S. banking supervisors' MRM guidance (SR 11-7) and similar directives from the OCC formalise this lifecycle approach, serving as key references for model governance in banking (The Federal Reserve, 2011).

2.1. The "Contained Failure" Assumption

In much of the finance AI literature—especially in areas like credit risk, AML/fraud analytics, market risk, and portfolio optimisation—there's a common assumption that model failures occur within a single firm. For example, a bank's credit model may drift, an asset manager's signals may degrade, a robo-advisor might misjudge risk tolerance, or an AML model could generate too many

false positives. Typical solutions proposed include improved validation, data governance, explainability, documentation, and accountability.

Recent regulatory developments support this approach. The PRA's SS1/23 encourages banks to adopt a strategic, enterprise-level model risk discipline, recognising model risk as a standalone category (Bank of England, 2024). The ECB's updated guide on internal models clarifies supervisory expectations and addresses machine-learning techniques, thereby enhancing governance, controls, and transparency. Similarly, the EBA focuses on the use of machine learning in IRB settings while emphasising compliant model deployment (Deloitte, 2023).

All of this is essential. However, it alone does not adequately address the evolving risks associated with AI in finance.

2.2. Interdependent Failure and System-Level Coupling

The literature often overlooks the shift from firm-level model risk to system-level coupling risk, where multiple institutions using similar AI infrastructures can lead to correlated behaviours and failures. The core issue isn't just "bad models," but rather shared dependencies and aligned decision-making processes: common foundation models, cloud services, data vendors, AI development platforms, and middleware APIs across banks and asset managers.

This situation creates model monoculture risk. Even if individual institutions meet their internal model risk management (MRM) standards, the overall financial system could still become fragile if many "independent" players essentially operate with the same AI frameworks. The Financial Stability Board (FSB) emphasises that vulnerabilities extend beyond firm governance to include third-party dependencies and market correlations linked to AI use, which can affect financial stability (FSB, 2025). Similarly, the Basel Committee notes that new technologies, including AI and cloud services, can heighten interconnectedness and create systemic risk.

In capital markets, the International Organisation of Securities Commissions (IOSCO) highlights the wide range of AI applications—from trading to advice, research, and surveillance—where feedback loops and synchronisation effects are likely to occur.

2.3. Why Existing MRM Lenses Do Not "See" Monoculture Risk

Traditional MRM frameworks are designed to address unique model failures. However, they are less effective at identifying four systemic pathways that can occur even when every organisation is "compliant."

1. Vendor concentration and correlated outages

When multiple firms depend on the same cloud or AI middleware, incidents such as operational failures, model updates, API limitations, or security issues can have widespread impacts, even if each firm has validated its own model.

2. Common model priors and shared blind spots

Fine-tuned models can inherit biases and flaws from their base models, leading to vulnerabilities across the sector that may seem independent at first glance.

3. Signal convergence and crowded positioning

When organisations absorb comparable alternative data and utilise similar analytical methods, their portfolios may converge on similar exposures. During times of stress, collective de-risking can intensify volatility and liquidity disruptions—a classic systemic phenomenon, now expedited by the rapid pace enabled by AI.

4. Compliance-driven convergence

Tighter regulatory expectations can, ironically, lead to greater monoculture. The pressure to comply may push organisations towards standardised, vendor-certified solutions and common control templates, which reduces diversity in model ecosystems, even as governance improves.

2.4. From “Model Quality” to “Cognitive Diversification”

The literature needs a clear connection between microprudential model risk and macro-level stability. Key questions include: What diversity exists in AI decision-making? What are the system’s critical points of failure? How quickly can institutions switch to alternative workflows if AI services decline? And how should supervisors assess adoption patterns when the risk is more about “too many institutions using the same intelligence” rather than “a flawed model”?

Model monoculture risk is a distinct issue and focuses on analysing the distribution of dependencies across the sector, aligning with emerging supervisory monitoring efforts.

3. Regulatory Tailwinds and Blind Spots: AI Act “High-Risk” Mapping, IOSCO Themes, and FSB Monitoring

Regulation is increasingly becoming a significant supportive force for AI governance in the finance sector; however, it might also create a blind spot by reinforcing compliance at the firm level while inadequately addressing system-wide connections and concentration. Thus, the developing policy framework (such as the EU AI Act, guidance for securities markets, and macroprudential oversight) should be interpreted as having dual effects: it can elevate baseline governance standards, but it may also inadvertently hasten technology convergence—precisely the scenario that heightens *model monoculture risk*.

3.1. AI Act: “High-Risk” Finance Mapping as Governance Tailwind

The EU AI Act classifies creditworthiness assessment and credit scoring as high-risk activities, imposing strict requirements for risk management, data governance, documentation, human oversight, and post-market monitoring. This approach encourages financial institutions to rethink how they develop and justify automated credit decisions, particularly at the intersection of statistical scoring, machine learning, and “automated decision-making” (Arnal, 2026; Hacker & Eber, 2025).

This is significant for banking and retail finance because credit decisions impact consumer protection, prudential risk, and discrimination law. The regulation clearly indicates that the future of AI in lending will be governed, documented, and auditable. In essence, AI governance is evolving from voluntary ethics to mandatory law (Arnal, 2026; Hacker & Eber, 2025).

Blind spot: A compliance-focused, high-risk environment can lead to standardisation among institutions. When facing similar obligations and strict oversight, organisations often turn to vendor-packaged tools and compliance-ready systems. This approach can enhance baseline controls, but it may also create shared dependencies, such as using the same providers and tools. The outcome is a paradox: improved model governance within firms but reduced cognitive diversity across the system.

3.2. IOSCO: Capital-Markets AI Risks That Look “Firm-Level” but Behave “System-Level”

In capital markets, IOSCO’s AI work outlines use cases such as trading, research, distribution/advice, and market surveillance. It highlights key concerns, including governance, transparency, data quality, and operational resilience, noting that AI affects investor outcomes and market functioning (IOSCO, 2025). Scholars responding to IOSCO’s consultation emphasise that the regulatory challenge lies not only in identifying internal risks within firms but also in understanding market-wide feedback loops. These loops involve similar signals, strategies, faster reactions, and competitive information races (Azzutti, 2025).

Blind spot: Capital markets supervision often fails to detect synchronisation risk until it manifests as crowding, volatility spikes, or liquidity stress. By then, AI will have already become deeply integrated. Research on AI and crises highlights this issue: AI can amplify vulnerabilities through shared information sources, rapid actions, and strategic complementarities, creating systemic risks rather than isolated model failures. The 2010 Flash Crash incident in the U.S. is one good example (Kirilenko, 2017).

3.3. FSB Monitoring Themes: The Macroprudential Bridge—and Its Data Gap Problem

The FSB is focusing on the link between microprudential governance and macrofinancial stability by monitoring AI adoption and its vulnerabilities. It highlights risks such as third-party dependencies and market correlations that lead to model monoculture (FSB, 2025). This perspective strengthens the argument that AI risks in finance extend beyond model explainability and validation to include sector-wide concentration and coupling (FSB, 2025).

Blind spot: Oversight is only as effective as the authorities' understanding of adoption trends. The FSB itself emphasises the real-world challenges of monitoring and the shortcomings in data, particularly the challenges of tracking AI usage integrated into third-party services, procurement processes, and workflow tools, rather than being identified as “a model” within a bank’s inventory.

3.4. Synthesis: Regulation Reduces “Bad AI,” but May Increase “Shared AI”

This suggests that:

- The AI Act enhances the regulatory standards for firms in high-stakes retail finance, particularly in credit and underwriting (Arnal, 2026; Hacker & Eber, 2025).
- IOSCO highlights how market-facing AI can create feedback loops and raise concerns about market integrity (IOSCO, 2025).
- The FSB prioritises adoption monitoring, third-party dependency, and correlation channels as key financial stability issues (FSB, 2025).

None of these regimes effectively addresses a key issue related to AI-driven economic disruption in finance (Rhee & Dogra, 2023). The risk of regulatory convergence is significant. Recent research shows that when regulators disclose imperfect information about banks’ financial health, it can lead to the misclassification of healthy banks as risky. This encourages banks to adopt portfolios that regulators consider “safe” (Rhee & Dogra, 2023). This dynamic creates portfolio convergence, described by this paper as model monoculture, which reduces diversification and increases system fragility. It argues that an optimal disclosure policy follows a non-linear “bang-bang” structure based on the severity of adverse selection (Rhee & Dogra, 2023). This insight applies to AI governance: when regulatory regimes define what constitutes compliant or “safe” AI practices, institutions may converge on similar models and data infrastructure, reducing diversity in the financial system.

How do we measure and govern the system’s exposure to shared intelligence substrates (providers, model families, data pipelines, middleware) before they lead to correlated failures?

This question leads into the next section, which explores the paper’s focus on the Model–Market–Middleware framework. It highlights the need for explicit measurement of “model monoculture risk” rather than dismissing it as a mere side-effect in traditional model risk management (Danielsson & Uthemann, 2025).

4. The M³ Framework (Model–Market–Middleware)

We distinguish between model risk at the firm level and monoculture risk at the system level, and we introduce the M³ Framework (Model–Market–Middleware) as a solution. This framework elucidates the progression of artificial intelligence in banking and investment management, transitioning from localised automation to systemic interconnection.

AI risk in finance stems not just from the models themselves, but from how those models interact with markets and the infrastructure that links them.

4.1. Overview: From Isolated Models to Interdependent Cognition

Traditional model risk frameworks mainly concentrate on:

- Data quality
- Validation processes
- Explainability
- Governance and oversight

These controls are essential. Nevertheless, they presume that models function within confined organisational boundaries.

The M³ Framework focuses on three interacting layers: model similarity, market synchronisation, and middleware concentration. Figure 1 illustrates how these layers combine to create system-level exposure that goes beyond typical firm-level model risk.

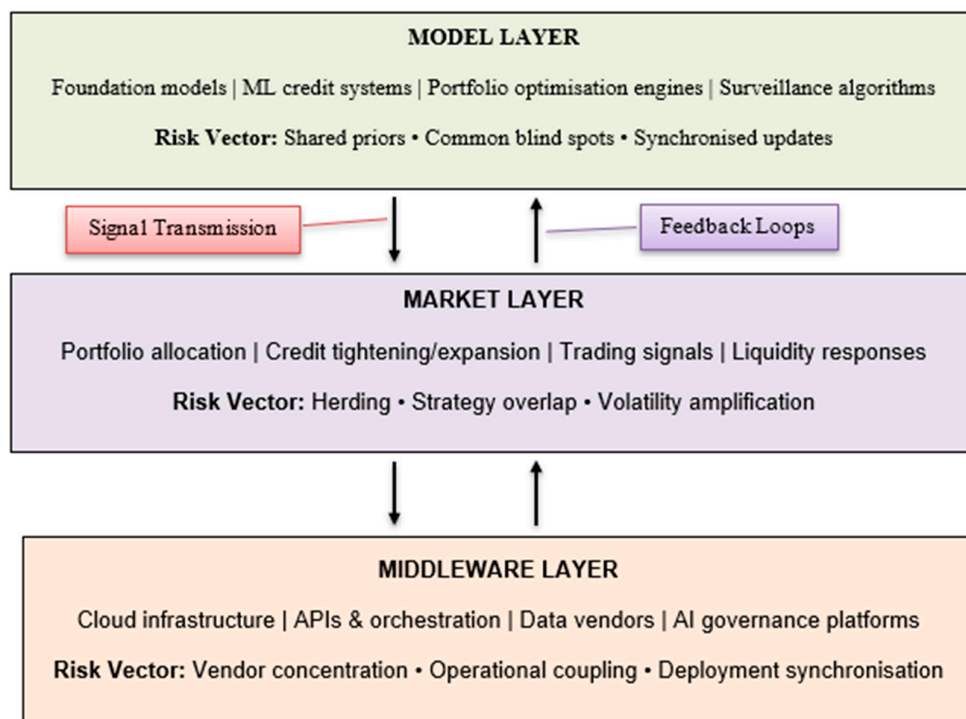


Figure 1. The M³ Framework: Model–Market–Middleware Coupling in AI-Driven Finance. The figure shows how shared AI models, synchronised market behaviours, and centralised middleware systems interact to create system-level exposure, known as "model monoculture risk." Arrows indicate feedback loops and dependencies among the layers. Source: Author's conceptualisation (2026).

4.2. The Model Layer: Shared Intelligence Substrates

The Model layer encompasses key AI architectures used in banking and investment management, including foundation models, machine learning systems, credit scoring engines, and portfolio optimisation tools. While these systems may seem unique to each institution due to tailored data and governance, there is a growing similarity in their underlying model families, training methods, and data dependencies. Many institutions depend on overlapping foundation model ecosystems and large-scale training architectures, as highlighted in recent literature (Bommasani et al., 2021). Even with customised fine-tuning, shared pre-training architectures can create structural similarities that impact outcomes.

This shared foundation is significant because it leads to common intelligence features among models. Institutions using different models may still share biases, limitations, and representational constraints from their underlying architectures. Research indicates that issues embedded during pre-training can persist in downstream adaptations, even with localised governance (Bender et al., 2021). If multiple banks and asset managers utilise models with similar architectures, they may respond to signals in similar ways, especially during stressful conditions. Changes in calibration or retraining at the upstream level can synchronise outputs across institutions that were previously considered independent.

The supervisory discourse on AI in finance has mainly focused on explainability, robustness, and model risk management, with a focus on validation, monitoring, and documentation (Basel Committee on Banking Supervision, 2024). While these issues are essential, they overlook the systemic risks associated with similar model structures across firms. Research has shown that correlated exposures and synchronised actions, rather than individual failures, are key drivers of systemic risk (Acharya, 2009). Thus, the main risk is not just model errors but rather model alignment, which reduces cognitive diversity in the financial system when similar architectures and training are used.

Indicators of model-layer monoculture include dependency on a few foundation model providers, shared external datasets, synchronised model updates, uniform benchmarking standards, and similar hyperparameter optimisation practices. When these factors align, it suggests that the financial sector tends to behave more like a network of interconnected systems than as independent entities. Consequently, the model layer is the first avenue through which AI shifts firm-level automation into broader systemic vulnerabilities.

4.3. *The Market Layer: Feedback Loops and Behavioural Synchronisation*

The Market layer examines how AI-generated outputs influence economic actions in banking and investment management. AI systems are increasingly used for asset allocation, credit decisions, liquidity management, trading execution, counterparty assessment, and risk limit calibration. At this level, the focus shifts from how models work internally to how their outputs impact market behaviour. When institutions use similar AI signals—such as in sentiment analysis or risk scoring—this can lead to synchronised economic responses.

Research on financial stability shows that systemic fragility often stems from correlated exposures and synchronised actions, rather than from individual failures (Acharya, 2009). AI can intensify this effect by speeding up and standardising signal interpretation. Studies in algorithmic trading indicate that widely adopted strategies can amplify volatility and decrease market resilience (Kirilenko et al., 2017; Menkveld, 2013). Although these findings were made before the widespread use of foundation models, the main idea remains: automation minimises behavioural variation and can speed up feedback loops under stress.

Recent research on algorithmic price formation highlights that algorithmic and AI-driven systems increasingly influence price discovery through complex, opaque processes. This creates "rational opacity," in which prices are efficient but their causes are unclear to human participants (Frimpong & Mamuti, 2026). This view aligns with the M³ framework, illustrating how algorithmic systems can limit interpretive diversity in markets before visible behavioural synchronisation occurs.

In AI-enhanced environments, the risk of convergence extends beyond high-frequency trading. Decisions regarding credit tightening, collateral adjustments, and capital redistribution may increasingly be influenced by AI-driven analytics integrated into internal dashboards and risk management systems. When these systems utilise similar data sources or employ analogous predictive frameworks, the risk-reduction thresholds may align across institutions. Research on procyclicality and amplification mechanisms indicates that such alignment can intensify downturns, especially when institutions make balance-sheet adjustments simultaneously (Brunnermeier & Oehmke, 2013).

Crucially, it is not necessary for AI to be identical across companies to create synchronisation. It is sufficient for models to be trained on overlapping datasets, respond to similar macroeconomic indicators, or achieve comparable performance metrics. During stressful periods, minor adjustments in signal calibration can trigger widespread portfolio rebalancing or liquidity withdrawals. In this way, AI heightens the responsiveness of collective behaviour: reactions become quicker, more automated, and less influenced by diverse managerial judgment.

Thus, indicators of Market-layer coupling include increased strategy overlap, convergence in factor exposures, synchronised adjustments to risk limits, clustering of AI-induced stop-loss mechanisms, and concentration in shared alternative data signals. When these indicators coincide

with Model-layer similarity and Middleware concentration, localised decision-making can lead to system-wide amplification. Therefore, the Market layer serves as the behavioural transmission medium through which AI-driven cognitive convergence turns into macrofinancial risk.

4.4. *The Middleware Layer: Infrastructure Concentration and Operational Coupling*

The Model layer captures shared intelligence architectures, while the Market layer reflects synchronised economic responses. The Middleware layer serves as the foundation for deploying, maintaining, and connecting AI systems. It encompasses cloud services, API-based model access, data vendors, orchestration platforms, compliance toolkits, model risk management software, and workflow integration systems. Unlike the visible Model layer, the Middleware layer often operates under the radar, embedded in procurement, DevOps pipelines, and vendor management.

Research on financial stability shows that concentration in critical financial infrastructure can lead to systemic vulnerabilities, even if individual institutions are well-managed. Duffie (2015) points out that central clearing infrastructures can become critical nodes, whose failures can ripple through the financial system. Although cloud providers and AI middleware are not clearinghouses, they pose similar risks: reliance on shared infrastructure raises concerns about systemic resilience.

Regulatory authorities have noted that increased digitalisation leads to greater dependence on third-party technology providers, raising worries about operational concentration and potential disruptions (Basel Committee on Banking Supervision, 2024). The banking sector's reliance on cloud service providers invites scrutiny, as outages can impact multiple institutions. This is considered operational risk, but in an AI-driven financial ecosystem, risks extend beyond availability. The concentration of middleware differs from operational risk in that it influences interpretive alignment rather than merely ensuring service continuity. Middleware synchronises deployment cycles, model updates, data ingestion pipelines, and compliance templates, fostering cognitive and procedural alignment across institutions.

The impact of this coupling effect becomes especially pronounced when AI services are utilised via standardised APIs or vendor-managed platforms. Refreshes of fundamental components, performance enhancements, retraining cycles, or security updates can ripple through client organisations in short timeframes. While this standardisation boosts efficiency and diminishes barriers to adoption, it also lessens diversity in deployment architectures. The outcome is not just operational interdependence but a synchronised evolution of AI systems across different companies.

Research on systemic risk highlights that fragility often stems from network structure and common vulnerabilities rather than merely from isolated weaknesses in balance sheets (Acemoglu, Ozdaglar, & Tahbaz-Salehi, 2015). A concentration in middleware creates a networked dependency framework in which technology providers become pivotal nodes. When this is coupled with similarities at the model layer and synchronisation at the market layer, this infrastructural concentration turns localised disturbances—like a model performance issue or service disruption—into widespread system disturbances.

Signs of middleware-layer monoculture include significant market share in the cloud, dependence on a single vendor for AI development platforms, execution of decisions via APIs embedded in core processes, uniform toolkits aimed at compliance, and a lack of diversity in failover options in AI deployment frameworks. Hence, the middleware layer serves as the structural conduit through which “independent” financial institutions become both operationally and cognitively linked. It transforms AI adoption from a firm-level capability into a collective, systemic foundation.

4.5. *Interaction Effects: When $M \times M \times M$ Becomes Systemic*

The analytical strength of the M^3 framework lies not in its separate layers but in their interconnections. Systemic fragility arises not simply from similarities among models, nor from market synchronisation alone, nor from middleware concentration in isolation. It manifests when these three factors align and reinforce each other. In these situations, artificial intelligence shifts from serving as a tool for firm-level optimisation to acting as a multiplier of correlated risk exposure.

Network-based systemic risk theory shows that shocks spread nonlinearly when interconnections amplify shared exposures (Acemoglu, Ozdaglar, & Tahbaz-Salehi, 2015). Within the M³ framework, the Model layer influences the interpretation of signals, the Market layer decides how those signals translate into economic actions, and the Middleware layer controls how quickly and uniformly those signals spread across institutions. When there is similarity at the model level, synchronisation at the market level, and concentration at the middleware level, the financial system becomes very responsive to disturbances from upstream sources.

Imagine a simplified scenario: an update to a foundation model alters the sentiment calibration in risk-assessment tools used by various asset managers. Because these institutions depend on comparable architectures (Model similarity), their portfolio dashboards display similar changes in perceived risk exposure. Automated allocation systems then trigger de-risking adjustments across linked positions (Market synchronisation). Given that deployment pipelines and APIs are commonly shared among institutions (Middleware concentration), these updates disseminate quickly and with minimal variation. As a result, individually rational adjustments lead to collective magnification.

The literature on financial stability has long highlighted that correlated deleveraging and shared exposures are key drivers of systemic crises (Brunnermeier & Oehmke, 2013; Acharya, 2009). The M³ interaction extends this framework to environments influenced by AI. The new aspect is not the presence of herding in markets—which is already well documented—but rather that AI might reduce the diversity of interpretations before herding occurs. Alignment is integrated earlier in the cognitive process.

Crucially, the interaction term is multiplicative instead of additive. A system with similar models but varied infrastructure may slow shock absorption. Conversely, a system with concentrated middleware but diverse models might still maintain interpretive variety. It is the alignment of the three layers—Model × Market × Middleware—that creates structural fragility. In such a scenario, localised calibration changes, vendor failures, or retraining periods can lead to synchronised behavioural reactions among institutions.

The consequence is that systemic risk associated with AI cannot be evaluated using a single dimension. Monitoring needs to assess cross-layer alignment: how similar are the foundational model architectures, how synchronised are market reactions, and how concentrated are the dependencies on infrastructure? Only by exploring the interaction space can regulators and risk managers identify the rise of model monoculture at the system-wide level.

5. The Model Monoculture Risk Index: Measuring System-Level AI Convergence

If the M³ Framework pinpoints the structural pathways through which the risks of AI-driven monoculture could arise, the subsequent inquiry is operational: how can financial entities and regulators recognise its accumulation before a correlated failure occurs? Conventional model risk inventories evaluate the quality of validation, documentation, and performance stability within an organisation. However, these instruments are not intended to assess cross-layer alignment or systemic convergence. The Model Monoculture Risk Index (MMRI) is suggested as a structured framework that transitions from qualitative to semi-quantitative evaluation of exposure across the Model, Market, and Middleware layers concurrently.

The MMRI is not intended to assess model accuracy. Instead, it gauges the decline in diversity within the AI decision-making framework in the finance system.

5.1. Conceptual Structure of the MMRI

The index is organised into three dimensions that correspond to the M³ Framework: Model Similarity (MS), Market Synchronisation (MK), and Middleware Concentration (MW). Every dimension represents a unique aspect of systemic alignment instead of overlapping operational characteristics.

1) Model Similarity Score (MS)

This dimension examines the extent of alignment in AI architectures and upstream dependencies among various institutions. Instead of focusing on the effectiveness of the models, it evaluates the shared structural features in the models' design and training fundamentals.

Key indicators include:

- Concentration within foundational model ecosystems.
- Shared pre-trained architectures or base model families
- Reliance on overlapping external datasets or feature libraries
- Synchronised update cycles or benchmarking standards

A higher MS score means institutions have a more uniform understanding of financial signals.

2) Market Synchronisation Score (MK)

This dimension evaluates how AI outputs influence economic actions by examining whether AI-assisted systems lead to consistent adjustments in portfolios, credit, or liquidity across firms.

Core indicators include:

- Overlap in strategy exposures or factor tilts
- Correlated signal thresholds triggering action
- Simultaneous recalibration of risk limits or collateral policies
- Concentration in common alternative data signals

A higher MK score means greater amplification potential due to synchronised behavioural responses.

3) Middleware Concentration Score (MW)

This dimension assesses the interconnection of infrastructure and the coordination of deployments. It examines how much institutions utilise common technological foundations that facilitate AI capabilities.

Primary indicators include:

- Cloud or AI platform market concentration
- Dependence on vendor-managed AI development environments
- API-based integration of AI into core decision workflows
- Limited fallback diversity in deployment architectures

A higher MW score indicates faster propagation speed and better systemic coupling via shared infrastructure.

5.2. The Multiplicative Logic

The MMRI is not a simple additive index because systemic exposure does not increase linearly with similarity, synchronisation, or concentration alone. Instead, fragility arises from interconnections across layers. Therefore, the framework uses a multiplicative structure to capture these interaction effects:

$$MMRI = f(MS \times MK \times MW)$$

The multiplicative structure describes a theoretical interaction instead of a fixed econometric model. Systemic exposure increases when model similarity, market synchronisation, and middleware concentration rise together; if any of these is low, the amplification is limited. Fragility arises from convergence across different layers, not from a single indicator. The function $f(\cdot)$ retains flexibility, enabling the MMRI to be implemented using qualitative scoring or nonlinear empirical models without relying on a specific estimation method.

5.3. Interpretive Zones

Table 1 outlines four structural exposure zones derived from different combinations of Model Similarity (MS), Market Synchronisation (MK), and Middleware Concentration (MW). These zones represent the relative levels of systemic alignment and vulnerability, rather than indicating crisis conditions.

Table 1. Structural Exposure Zones of the Model Monoculture Risk Index (MMRI).

Zone	Structural Condition	Systemic Exposure
I	Low MS, Low MK, Low MW	High diversity, low systemic AI coupling
II	High MS only	Cognitive similarity but behavioural dispersion
III	High MS + High MK	Amplification risk without infrastructure fragility
IV	High MS + High MK + High MW	Full monoculture exposure

Source: Author's conceptualisation (2026).

Zone IV is characterised by a structural arrangement closely linked to systemic fragility driven by AI, in which factors such as similarity, synchronisation, and infrastructure concentration amplify one another. The shift between zones can occur gradually, underscoring the importance of ongoing observation rather than simple binary categorisation.

5.4. Supervisory and Strategic Implications

The MMRI framework impacts three different levels:

For regulators: AI adoption monitoring should focus not just on governance compliance, but also on metrics related to structural concentration. Data collection should include mapping vendor reliance and analysing inter-institution signal dependencies.

For chief risk officers: AI diversification serves as a strategic tool for resilience. Institutions can intentionally introduce variety in model types, data pipelines, or cloud services to minimise systemic alignment.

For boards and policymakers: Enhancements in AI efficiency should be weighed alongside the importance of maintaining systemic diversity. As capital adequacy safeguards against balance-sheet vulnerabilities, cognitive diversity could serve as a defence against shocks to correlations induced by AI.

5.5. Contribution of the MMRI

The MMRI contributes to the literature in three key aspects. It shifts AI risk evaluation from model-level validation to system-level architecture. Then, it combines both microprudential and macroprudential viewpoints. Lastly, it presents quantifiable indicators for a risk category that is often discussed but infrequently put into practice. In doing so, it redefines AI governance in finance as not merely a matter of compliance but a challenge of structural diversification.

6. Governance Implications and Future Organisational Models

The rise of model monoculture risk calls into question the common belief that stronger organisational-level governance of AI inherently improves resilience across the entire system. Although compliance measures like the AI Act and risk management standards enhance transparency, documentation, and accountability, they do not guarantee the maintenance of diversity within the financial ecosystem. The M³ framework indicates that governance needs to progress from merely validation and explainability to encompass structural diversification and infrastructural resilience.

6.1. From Model Governance to Architecture Governance

Traditional AI governance frameworks emphasise managing the model lifecycle, including development, validation, monitoring, documentation, and human oversight. While these controls are

crucial, they function mainly within the confines of institutions. The systemic dynamics highlighted in this paper suggest the necessity for an additional layer of architecture governance—one that is supervisory and strategic, assessing concentration, interdependence, and alignment across different layers. Ongoing supervisory conversations about digitalisation and risks associated with third-party concentration hint at this broader structural viewpoint (Basel Committee on Banking Supervision, 2024).

Architecture governance necessitates institutions and regulators to oversee:

- Concentration in foundation model ecosystems
- Overlapping reliance on cloud and middleware providers
- Correlation in AI-triggered decision thresholds
- Structural fallback capacity and model diversity

This approach builds on operational resilience oversight by applying it to the cognitive domain. In AI-driven finance, resilience means not just keeping systems running, but also maintaining diverse interpretations across institutions.

6.2. *The Strategic Case for Cognitive Diversification*

From a managerial view, adopting AI is often seen as a quest for greater efficiency and predictive capability. The pressure to compete drives organisations toward the most effective and compliant solutions. However, systemic risk theory shows that similar vulnerabilities and coordinated actions can increase fragility, even if each organisation acts rationally (Acharya, 2009). As a result, too much technological convergence may undermine overall resilience.

Strategic leaders in banking and asset management may face a new trade-off:

Efficiency versus cognitive diversification.

Similar to how diversifying a portfolio lowers the risk associated with correlated financial issues and network effects (Acemoglu, Ozdaglar, & Tahbaz-Salehi, 2015), spreading out across various AI architectures, data sources, and deployment infrastructures could diminish the risk of correlated cognitive failures. Organisations might intentionally choose to sustain:

- Multiple model families
- Diverse data ingestion pipelines
- Hybrid cloud strategies
- Independent validation environments

This kind of diversification might seem inefficient on a small scale, but it tends to provide stability on a larger scale. In this way, AI governance evolves from merely a risk management role into a challenge of strategic resilience design.

6.3. *Supervisory Evolution: Monitoring Alignment, Not Just Accuracy*

For regulators and macroprudential authorities, the implications of policy are nuanced yet significant. Current supervisory frameworks focus on evaluating whether models are precise, equitable, well-documented, and understandable. The M³ analysis indicates that regulators should also examine the similarities of these models across different institutions and the extent to which their outputs align during times of stress. Recent discussions on financial stability regarding AI implementation highlight the need to oversee dependencies on third parties and systemic interconnections (Financial Stability Board, 2024).

Monitoring might involve:

- Sector-wide vendor concentration mapping
- Cross-institution signal dependency analysis
- AI update synchronisation reporting
- Cloud concentration stress testing
- MMRI-style heatmap tracking

This type of monitoring does not require regulators to closely examine proprietary algorithms. Instead, it calls for insight into dependency structures and inter-layer connections. This signifies a transition from overseeing models to overseeing architectures.

6.4. Future Organisational Models: From AI Integration to AI Dependency

Artificial intelligence is changing financial products and how institutions are designed. Research on digital transformation indicates that AI and data-driven systems are altering decision-making authority, organisational coordination, and managerial roles (Brynjolfsson & McAfee, 2014). As AI increasingly mediates decisions, organisations may shift from owning models to relying on infrastructure within larger technological ecosystems. Middleware providers and foundation model ecosystems are becoming integral to governance, influencing institutional capabilities in indirect ways.

Future financial organisations might therefore be defined by:

- Distributed cognition in human–AI interfaces
- Vendor-mediated intelligence layers
- Continuous model updating cycles
- Embedded compliance automation

In contemporary organisational environments, leadership roles are undergoing significant transformation. Chief Risk Officers and Chief Information Officers are increasingly required to collaborate closely with procurement, strategy, and board-level oversight functions to assess and address systemic alignment risks effectively. Furthermore, the governance of artificial intelligence is evolving into a responsibility that spans the entire enterprise, moving beyond a narrow focus on technical or compliance aspects.

6.5. Broader Economic Implications

AI's impact on finance goes beyond productivity and job displacement; it reshapes information processing and the formation of expectations. According to network-based systemic risk theory, fragility arises from common vulnerabilities and interdependence rather than isolated issues (Acemoglu et al., 2015). When AI alters risk perception across institutions, it affects credit cycles, asset pricing, and liquidity.

Maintaining diversity—analytical, infrastructural, and behavioural—can serve as a stabilising factor in AI-driven financial systems. Just as capital buffers safeguard against balance-sheet weaknesses, architectural diversity can protect against collective cognitive vulnerabilities.

6.6. Concluding Perspective

Artificial intelligence offers significant improvements in analysis and efficiency in finance. However, the evidence indicates that it is in how it changes systemic interdependence. The M³ Framework shows that model similarity, market synchronisation, and middleware concentration combine to create a risk of model monoculture that traditional governance cannot address.

The challenge is not to limit AI adoption but to ensure diversity and resilience within its structure. Financial stability in the AI era will rely on more than just capital adequacy and liquidity; it will also require maintaining diversity in market decision-making processes.

Ultimately, AI governance in finance needs to focus on creating systems that stay diverse, even as intelligence is shared.

7. Conclusion

Artificial intelligence is not just enhancing finance; it is reshaping it. As banks and financial markets incorporate foundational models, shared data systems, and vendor-driven AI ecosystems, the focus of risk transitions from individual model failures to systemic convergence. The main

argument of this paper is simple: when institutions think and act alike while utilising the same technological framework, vulnerability can arise even without any clear mistakes.

The idea of model monoculture risk reinterprets AI integration as a structural shift in the financial system. Employing the M³ Framework—Model, Market, and Middleware—this paper has illustrated how uniformity in AI frameworks, synchronisation in economic responses, and concentration in infrastructural dependencies interact multiplicatively. Individually rational optimisation can, in conditions of convergence, lead to collectively heightened exposure. In AI-driven environments, diversity is not merely a competitive advantage; it serves as a stabilising mechanism.

The suggested Model Monoculture Risk Index (MMRI) shifts the governance perspective from validation to alignment. The focus is no longer solely on whether models are accurate and understandable but on whether they are becoming structurally similar across various institutions. Maintaining financial stability in the age of AI may therefore rely less on refining individual algorithms and more on upholding heterogeneity within the entire system.

The goal is not to hinder innovation but to ensure that efficiency does not quietly erode diversity. In financial markets increasingly mediated by shared intelligence systems, resilience may depend on maintaining sufficient variation to withstand common shocks. AI has the potential to create more intelligent markets. The question remains whether these markets will remain sufficiently diverse to maintain stability, a decision influenced by governance.

References

- Acemoglu, B.D., Ozdaglar, A. & Tahbaz-Salehi, A. (2015). Systemic Risk and Stability in Financial Networks. *American Economic Review*, 105, 564–608. <https://doi.org/10.1257/aer.20130456>
- Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability*, 5(3), 224–255. <https://doi.org/10.1016/j.jfs.2009.02.001>
- Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., & Stein, M. (2025). Intelligent financial system: How AI is transforming finance. *Journal of Financial Stability*, 101472. <https://doi.org/10.1016/j.jfs.2025.101472>
- Arnal, J. (2026). Unfit for purpose? The legal maze of credit scoring under EU law. *Journal of Banking Regulation*, 27(2). <https://doi.org/10.1057/s41261-026-00311-7>
- Azzutti, A. (2025). Artificial Intelligence in Capital Markets: A Legal Scholar's Response to the IOSCO Consultation Report. Available at SSRN: <https://ssrn.com/abstract=5213493> or <http://dx.doi.org/10.2139/ssrn.5213493>
- Bank of England (2024). SS1/23 – Model risk management principles for banks. [www.bankofengland.co.uk](https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss). <https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss>
- Basel Committee on Banking Supervision (2024) Digitalisation of finance. <https://www.bis.org/bcbs/publ/d575.pdf>
- Bender, E. M., Geburu, T., McMillan-Major, A., & Shmittchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT'21)* (pp. 610–623). Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>
- BIS (2025, June 26) Financial stability implications of artificial intelligence - Executive Summary. [Bis.org](https://www.bis.org/fsi/fsisummaries/exsum_23904.htm). https://www.bis.org/fsi/fsisummaries/exsum_23904.htm
- Boeddu, G., E. Feyen, S. Martínez Jaramillo, S. Mesquita, Y. Palta, A. Sarkar, S. Sinha y A. Gutiérrez Traverso (2025). AI for financial sector supervision: New evidence from emerging market and developing economies. CEPR. <https://cepr.org/voxeu/columns/ai-financial-sector-supervision-new-evidence-emerging-market-and-developing-economies>

- Bommasani R., Hudson, D. A., Ehsan Adeli, Altman, R. B., Arora, S., Sydney von Arx, Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N. S., Chen, A. T., Creel, K., Davis, J., Demszky, D., & Donahue, C. (2021). On the Opportunities and Risks of Foundation Models. ArXiv (Cornell University). <https://doi.org/10.48550/arxiv.2108.07258>
- Brunnermeier, M.K. & Oehmke, M. (2012). The Maturity Rat Race. *The Journal of Finance*, 68, 483–521. <https://doi.org/10.1111/jofi.12005>
- Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company.
- Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024). Regulating AI in the Financial Sector: Recent Developments and Main Challenges. Bank for International Settlements FSI Insights on Policy Implementation (No. 63).
- Danielsson, J., & Uthemann, A. (2025). Artificial intelligence and financial crises. *Journal of Financial Stability*, 80, 101453. <https://doi.org/10.1016/j.jfs.2025.101453>
- Deloitte (2023, September 19), Adoption of advanced machine learning techniques for IRB Models. Deloitte. <https://www.deloitte.com/dk/en/services/risk-advisory/research/adoption-of-advanced-machine-learning-techniques-for-irb-models.html>
- Duffie, D. (2024). Resolution of Failing Central Counterparties. Retrieved February 24, 2026, from <https://www.hoover.org/sites/default/files/research/docs/makingfailurefeasible-ch4.pdf>
- EBA (2020), The EBA supports the Commission's proposal for a new Digital Finance Strategy for Europe | European Banking Authority. Europa.eu. <https://www.eba.europa.eu/publications-and-media/press-releases/eba-supports-commissions-proposal-new-digital-finance>
- Financial Stability Board. (2024). The Financial Stability Implications of Artificial Intelligence. <https://www.fsb.org/uploads/P14112024.pdf>
- Foy, S. (2025, December 31). AI forecast to put 200,000 European banking jobs at risk by 2030. *Financial Times*. <https://www.ft.com/content/71e12f85-1edb-4156-8cb5-3fe8aef36d93>
- Frimpong, V., & Mamuti, A. (2026). Opaque Price Control and Algorithmic Authority in Financial Markets. *Encyclopedia*, 6(1), 19. <https://doi.org/10.3390/encyclopedia6010019>
- FSB (2025, October 10). Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector. Financial Stability Board. <https://www.fsb.org/2025/10/monitoring-adoption-of-artificial-intelligence-and-related-vulnerabilities-in-the-financial-sector/>
- Hacker, P., & Eber, M. (2025). The Future of Credit Underwriting and Insurance Under the EU AI Act: Implications for Europe and Beyond. *Harvard Data Science Review*, 7(3). <https://doi.org/10.1162/99608f92.171157d2>
- IOSCO (2025) Artificial Intelligence In Capital Markets: Use Cases, Risks, And Challenges. The Board of The International Organisation of Securities Commissions Consultation Report. <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf>
- Kholia, T. (March 13, 2025). Artificial Intelligence in Capital Markets: Use Cases, Risks, and Regulatory Challenges. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.5179985>
- Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The Flash Crash: High-Frequency Trading in an Electronic Market. *The Journal of Finance*, 72, 967–998. <https://doi.org/10.1111/jofi.12498>
- Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712–740. <https://doi.org/10.1016/j.finmar.2013.06.006>
- Perez-Cruz, F., Jermy Prenio, R., F., & Yong, J. (2025, September 8). Managing explanations: how regulators can address AI explainability. *Bis.org*. <https://www.bis.org/fsi/fsipapers24.htm>

Rhee, K. & Dogra, K. (2023). Stress Tests and Model Monoculture. *Journal of Financial Economics (JFE)*, Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3208233>

The Federal Reserve (April 4, 2011). The Fed - Supervisory Letter SR 11-7 on guidance on Model Risk Management. [federalreserve.gov](https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm). <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>

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.