



## The Compliance Horizon: Anticipating Regulatory Change in Financial Services and Artificial Intelligence

Neha Tyagi

Senior Vice President, Bank of New York, USA

**ABSTRACT:** The convergence of artificial intelligence (AI) and financial services is precipitating a fundamental reconfiguration of the global compliance landscape. This paper examines the emergent regulatory frameworks being developed across major jurisdictions — including the European Union's AI Act, the United States' evolving federal approach, and the United Kingdom's principles-based regime — and analyses their implications for financial institutions navigating an era of algorithmic decision-making, automated risk management, and generative AI deployment.

Drawing on regulatory text analysis, institutional theory, and a comparative policy framework, we argue that traditional compliance architectures are inadequate to address the velocity and opacity of AI-driven financial processes. We identify three structural tensions in contemporary regulation: the innovation-stability paradox, the explainability imperative, and the jurisdictional fragmentation problem. The paper proposes a forward-looking 'Dynamic Compliance Architecture' (DCA) model, offering practitioners and policymakers a structured methodology for anticipating, absorbing, and adapting to regulatory change in real time.

Our findings have significant implications for Chief Compliance Officers, RegTech vendors, central banks, and international standard-setting bodies. We conclude that proactive regulatory horizon-scanning, embedded in institutional governance frameworks, is no longer optional — it is a strategic imperative.

**KEYWORDS:** Regulatory compliance, artificial intelligence, financial services, EU AI Act, RegTech, model risk, algorithmic accountability, dynamic compliance architecture

### I. INTRODUCTION

The financial services industry has, across successive decades, demonstrated a remarkable capacity for regulatory adaptation. From the post-2008 Basel III capital accords to the sweeping conduct reforms embodied in MiFID II, institutions have repeatedly retooled their compliance infrastructures in response to externally imposed frameworks. Yet the current moment presents a qualitatively different challenge: the integration of artificial intelligence into core financial operations is generating regulatory complexity at a pace and scale that outstrips conventional compliance methodologies (Arner, Barberis & Buckley, 2017).

AI systems — encompassing machine learning (ML) models, large language models (LLMs), and automated decision engines — are now embedded across the financial services value chain. Credit underwriting, fraud detection, algorithmic trading, customer onboarding, and AML/KYC processes are increasingly AI-mediated (Financial Stability Board, 2022). This proliferation has not been lost on regulators. The period from 2022 to 2025 has witnessed an unprecedented volume of regulatory output targeting AI in finance, from the European Union's landmark AI Act (2024) to guidance from the Basel Committee on Banking Supervision (BCBS), the Financial Conduct Authority (FCA), and the US Office of the Comptroller of the Currency (OCC).

This paper makes three contributions. First, it provides a comparative analysis of the principal regulatory frameworks governing AI in financial services across major jurisdictions. Second, it identifies the structural tensions that characterise contemporary AI compliance. Third, it introduces the Dynamic Compliance Architecture (DCA) — a practitioner-oriented model designed to help institutions move from reactive compliance to anticipatory governance.



The remainder of the paper is structured as follows: Section 2 reviews the relevant literature. Section 3 provides the comparative regulatory analysis. Section 4 identifies structural tensions. Section 5 introduces the DCA framework. Section 6 discusses implications for practice and policy. Section 7 concludes.

## II. LITERATURE REVIEW

The regulatory governance of financial technology has attracted substantial academic attention since the post-GFC reforms. Zetzsche et al. (2017) established a foundational taxonomy of FinTech regulation, distinguishing between entity-based, activity-based, and outcomes-based approaches — a framework that retains analytical relevance as AI applications proliferate. Notably, they warned against regulatory frameworks that conflate technological novelty with systemic risk, a caution that resonates acutely in the AI context.

The concept of 'regulatory lag' — the structural delay between technological innovation and commensurate regulatory response — has been extensively theorised (Zetzsche, Buckley & Arner, 2021). In the AI domain, this lag is exacerbated by the opacity of machine learning systems, which present fundamental challenges to transparency-based regulatory paradigms (Wachter, Mittelstadt & Russell, 2018). The 'black box' problem, as it has come to be known, creates specific difficulties for regulators who require auditability and interpretability as preconditions of accountability.

Model risk management has emerged as a critical compliance subdiscipline. The US Federal Reserve's SR 11-7 guidance (2011) and its UK equivalent from the Prudential Regulation Authority (PRA) have framed model risk as a function of model error, misapplication, and inadequate governance. However, as Brundage et al. (2018) observe, AI systems introduce additional dimensions of risk — including distributional shift, adversarial vulnerability, and emergent behaviour — that are not adequately captured by pre-AI model risk frameworks.

From an institutional theory perspective, DiMaggio and Powell's (1983) concept of isomorphism offers a valuable lens for understanding regulatory convergence and divergence. As financial institutions respond to AI regulation, they are subject to coercive pressures (from regulators), mimetic pressures (from peer institutions), and normative pressures (from professional bodies and industry associations). The resulting compliance responses may not always be optimal from a risk-management standpoint.

The RegTech literature — emerging as a distinct field from approximately 2015 — has increasingly addressed the role of technology in compliance automation (Colaert, 2018; Anagnostopoulos, 2018). RegTech solutions promise to resolve some of the scalability challenges posed by complex AI regulation, but they also introduce secondary compliance risks, including vendor dependency, algorithmic auditing challenges, and questions of regulatory acceptance (Arner, Barberis & Buckley, 2017).

Most recently, the generative AI wave — epitomised by large language models such as GPT-4, Gemini, and Claude — has introduced new compliance considerations around hallucination risk, intellectual property, consumer protection, and systemic concentration (Bommasani et al., 2021). The integration of generative AI into compliance functions themselves — for regulatory change management, contract review, and risk reporting — adds a further recursive dimension to the compliance challenge.

## III. COMPARATIVE REGULATORY LANDSCAPE

### 3.1 The European Union: A Risk-Based Framework

The EU AI Act, formally adopted in March 2024, represents the world's first comprehensive horizontal regulatory framework for AI. Its risk-stratified architecture classifies AI systems into four tiers: unacceptable risk (prohibited), high risk (subject to conformity assessments), limited risk (transparency obligations), and minimal risk (self-regulatory). For financial services, the high-risk category is particularly significant, encompassing AI used in creditworthiness assessment, insurance risk evaluation, and AML screening (European Parliament, 2024).

High-risk financial AI systems under the AI Act must satisfy requirements across six dimensions: data governance, technical documentation, transparency, human oversight, accuracy and robustness, and cybersecurity. The Act's territorial scope — applying to AI systems deployed in the EU regardless of where the developer is based — gives it



extraterritorial significance comparable to the GDPR. Compliance obligations take effect on a phased basis, with general-purpose AI model obligations applying from August 2025.

The Act intersects with existing financial regulation in complex ways. The European Banking Authority (EBA) has issued guidance on the interplay between AI Act obligations and those under CRR/CRD, MiFID II, and DORA — the Digital Operational Resilience Act — which itself imposes significant requirements on financial entities' ICT risk management (EBA, 2023). The resulting regulatory matrix is formidable, requiring institutions to navigate concurrent obligations with overlapping but non-identical scope.

### 3.2 The United Kingdom: Principles-Based Pragmatism

Post-Brexit, the UK has adopted a distinctively different regulatory philosophy. Rather than sector-agnostic horizontal legislation, the UK's AI regulatory approach — articulated in the 2023 AI White Paper and subsequent FCA and PRA guidance — assigns responsibility for AI governance to existing sector regulators using a principles-based framework centred on five cross-sectoral principles: safety and security; transparency and explainability; fairness; accountability and governance; and contestability and redress (DSIT, 2023).

The FCA's approach to AI in financial services has been characterised by a 'test and learn' orientation, facilitated through the Regulatory Sandbox and the AI Lab. The FCA's Consumer Duty (PS22/9), which came into force in July 2023, has significant AI implications, requiring firms to demonstrate that AI-driven product recommendations and customer communications deliver good outcomes for retail consumers (FCA, 2022). The PRA has complemented this with its Supervisory Statement SS1/23 on model risk management, which explicitly extends to AI models and establishes expectations around validation, governance, and documentation (PRA, 2023).

### 3.3 The United States: A Fragmented Federal Landscape

The United States lacks a federal AI Act equivalent, producing a fragmented landscape in which financial AI is regulated through a patchwork of agency guidance, state legislation, and existing statutory frameworks. The Biden Administration's Executive Order on AI (October 2023) established broad principles and directed federal agencies to develop sector-specific implementation plans, but stopped short of binding AI-specific legislation (White House, 2023).

For financial institutions, the relevant regulatory corpus includes: OCC guidance on model risk management; the CFPB's focus on algorithmic discrimination in credit decisions; FINRA's rules on algorithmic trading; and FinCEN's emerging expectations for AI in AML systems. The SEC has signalled heightened scrutiny of AI-related disclosures and investment adviser use of predictive analytics (SEC, 2023). The Federal Reserve's SR 11-7 remains the bedrock of model risk governance, though its 2011 vintage predates modern AI by over a decade, creating interpretive challenges.

### 3.4 Emerging Markets and International Standards

Beyond the three major jurisdictions, a number of emerging market regulators have developed notable AI governance frameworks. The Monetary Authority of Singapore's (MAS) FEAT principles — Fairness, Ethics, Accountability, and Transparency — established a model for AI governance in finance that has influenced jurisdictions across Southeast Asia (MAS, 2019). The Reserve Bank of India and the Central Bank of Brazil have both issued consultation papers on AI risk management in 2024, reflecting growing regulatory momentum in the Global South.

At the international level, the Financial Stability Board (FSB) has published reports on AI/ML in financial services that highlight cross-border regulatory coordination challenges (FSB, 2022). The Basel Committee on Banking Supervision has addressed AI in the context of model risk, operational risk, and credit risk, while IOSCO has examined AI use by asset managers and investment firms. These international standards serve as important soft-law inputs for jurisdictions developing domestic frameworks.

## IV. STRUCTURAL TENSIONS IN AI COMPLIANCE

### 4.1 The Innovation-Stability Paradox

A foundational tension in AI financial regulation is the paradox between innovation facilitation and systemic stability. Regulatory intervention that is too prescriptive risks suppressing beneficial AI adoption and entrenching incumbent institutions; intervention that is too permissive risks the accumulation of systemic risk through opaque, correlated, and potentially destabilising AI-driven behaviours (Philippon, 2019). The EU AI Act's risk-stratification approach



represents one attempt to navigate this paradox, but its administrative burden has drawn criticism from FinTech innovators (European Parliament, 2024).

The herding dynamic presents a specific concern: as institutions converge on similar AI models and training datasets — particularly in credit risk and fraud detection — the financial system may exhibit increased homogeneity, reducing the diversity of risk-taking behaviour that underpins systemic resilience (FSB, 2022). This is a regulatory externality that firm-level compliance frameworks are ill-equipped to address.

## 4.2 The Explainability Imperative

Regulatory frameworks across jurisdictions increasingly demand explainability as a precondition of AI deployment in consequential financial decisions. The EU's GDPR (Article 22) established a right to explanation for automated decisions; the AI Act reinforces this with transparency requirements for high-risk systems; and the FCA's Consumer Duty mandates that firms can demonstrate how outcomes have been achieved. Yet explainability remains technically contested (Wachter, Mittelstadt & Russell, 2018).

The distinction between interpretability (the degree to which a human can understand a model's mechanics) and explainability (the degree to which an explanation of a specific decision can be provided post-hoc) is frequently elided in regulatory guidance. This creates compliance uncertainty: an institution may invest significantly in SHAP-based explanation tooling only to find that a regulator's conception of explainability requires something qualitatively different. The emergent field of 'regulatory explainability' — calibrated to the specific requirements of financial supervisors — represents a critical area for future research and standardisation (Adadi & Berrada, 2018).

## 4.3 The Jurisdictional Fragmentation Problem

Global financial institutions operating across multiple jurisdictions face a compliance matrix of formidable complexity. An AI model used in credit underwriting must simultaneously satisfy the EU AI Act's conformity assessment requirements, the UK PRA's SS1/23 validation expectations, the OCC's model risk management guidance, and MAS's FEAT principles — with no guarantee of mutual recognition or regulatory equivalence (Arner, Barberis & Buckley, 2017). The compliance cost of this fragmentation is substantial and disproportionately burdens smaller, cross-border institutions.

The prospect of regulatory arbitrage — institutions structuring AI development and deployment to minimise compliance obligations — is a related concern. The FSB has identified jurisdiction-shopping as an emerging risk in AI governance, particularly for cloud-based AI infrastructure that can be readily relocated across borders (FSB, 2022). Without greater international coordination, fragmentation is likely to intensify as domestic regulatory activity accelerates.

## V. THE DYNAMIC COMPLIANCE ARCHITECTURE (DCA) FRAMEWORK

In response to the structural tensions identified above, this paper introduces the Dynamic Compliance Architecture (DCA), a governance model designed to enable financial institutions to anticipate and adapt to regulatory change rather than merely react to it. The DCA is built on four interlocking pillars:

1. **Regulatory Horizon Scanning (RHS):** A systematic, institutionalised process for monitoring regulatory developments across relevant jurisdictions, international standards bodies, and supervisory communications. RHS functions should be supported by AI-driven monitoring tools capable of synthesising high volumes of regulatory output.
2. **Adaptive Policy Mapping (APM):** A methodology for mapping regulatory requirements onto institutional policies, processes, and AI systems in real time. APM requires a living repository of regulatory obligations — a 'compliance ontology' — that can be updated as new requirements emerge and cascaded through the organisation.
3. **Embedded Governance Controls (EGC):** The integration of compliance controls at the point of AI system design and development, rather than as ex-post overlays. This 'compliance-by-design' approach operationalises the principle of privacy-by-design (GDPR, Art. 25) at the AI governance level, embedding accountability, auditability, and explainability requirements into the AI development lifecycle.
4. **Continuous Assurance Loops (CAL):** A continuous, automated testing and assurance regime that monitors AI model performance, bias metrics, explainability outputs, and regulatory compliance indicators in real time — replacing the periodic model validation cycle with an always-on supervisory capability.



The DCA framework is intended to operate as an integrated system rather than a collection of discrete initiatives. Its effectiveness depends on senior leadership commitment, cross-functional collaboration between legal, compliance, risk, and technology functions, and meaningful investment in RegTech infrastructure. Critically, the DCA requires the institutionalisation of regulatory expertise within AI development teams — what might be termed a 'compliance engineer' capability — rather than relying exclusively on post-hoc compliance review.

## VI. IMPLICATIONS FOR PRACTICE AND POLICY

### 6.1 Implications for Financial Institutions

The DCA framework has immediate practical implications for Chief Compliance Officers and Chief Risk Officers. The most pressing is the need to restructure compliance functions to engage proactively with the AI development lifecycle. Compliance must evolve from a control function — checking finished outputs for regulatory alignment — to a design function embedded in AI governance from ideation through deployment and ongoing monitoring.

Institutions should also invest in regulatory intelligence capabilities. The volume and complexity of AI-related regulatory output across major jurisdictions now exceeds what traditional legal and compliance teams can absorb through manual processes. AI-powered regulatory change management tools — capable of ingesting, classifying, and mapping regulatory documents at scale — are becoming a compliance necessity rather than a luxury (Colaert, 2018).

### 6.2 Implications for Regulators and Policymakers

For regulators, the analysis suggests that principles-based frameworks — which allow institutions flexibility in implementation — may be better suited to the pace and diversity of AI innovation than prescriptive rule-based approaches. However, principles-based regulation is only effective where supervisors have the technical expertise to assess compliance meaningfully. Building AI literacy within regulatory bodies is an urgent priority (DSIT, 2023).

International coordination represents the most significant policy gap identified in this paper. The FSB, Basel Committee, and IOSCO are well-placed to develop common international standards for financial AI governance, reducing the compliance costs of jurisdictional fragmentation and limiting opportunities for regulatory arbitrage. The development of mutual recognition frameworks — analogous to those in derivatives regulation — warrants serious consideration.

## VII. CONCLUSION

The compliance horizon in financial services AI is simultaneously crowded and uncertain. Regulatory activity across major jurisdictions is intensifying, yet the frameworks being developed are frequently inconsistent, technically contested, and administratively burdensome. For financial institutions, the challenge is not simply to comply with today's rules, but to build the organisational capability to anticipate, interpret, and adapt to tomorrow's.

This paper has argued that traditional reactive compliance models are inadequate for the AI era, and has proposed the Dynamic Compliance Architecture as a structural response. The DCA's four pillars — regulatory horizon scanning, adaptive policy mapping, embedded governance controls, and continuous assurance loops — represent a practical pathway toward anticipatory compliance governance.

The stakes are significant. AI-driven financial systems that fail to meet evolving regulatory standards represent not only legal and reputational risk for individual institutions, but potential systemic risk for the broader financial system. Equally, regulatory frameworks that suppress beneficial AI innovation impose their own social costs. Navigating this tension — with rigour, intelligence, and foresight — is among the defining challenges of contemporary financial governance.

Future research should empirically test the DCA framework in institutional settings, examine the effectiveness of different regulatory approaches in managing AI-related financial risk, and explore the design of international coordination mechanisms for AI governance in finance. The compliance horizon is expansive; the need to map it has never been more urgent.



## REFERENCES

1. Adadi, A. & Berrada, M. (2018). Peeking inside the black box: A survey on explainable artificial intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
2. Anagnostopoulos, I. (2018). FinTech and RegTech: Impact on regulators and banks. *Journal of Economics and Business*, 100, 7–25.
3. Arner, D.W., Barberis, J. & Buckley, R.P. (2017). FinTech, RegTech, and the reconceptualization of financial regulation. *Northwestern Journal of International Law & Business*, 37(3), 371–413.
4. Basel Committee on Banking Supervision (BCBS). (2022). Principles for the Sound Management of Operational Risk. Bank for International Settlements.
5. Bommasani, R., Hudson, D.A., Adeli, E. et al. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.
6. Brundage, M., Avin, S., Clark, J. et al. (2018). The malicious use of artificial intelligence: Forecasting, prevention and mitigation. Future of Humanity Institute, University of Oxford.
7. Colaert, V. (2018). RegTech as a response to regulatory expansion in the financial sector. *KU Leuven Legal Studies Research Paper*, 2018/6.
8. Department for Science, Innovation and Technology (DSIT). (2023). A Pro-Innovation Approach to AI Regulation. UK Government White Paper. HMSO.
9. DiMaggio, P.J. & Powell, W.W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160.
10. European Banking Authority (EBA). (2023). Report on Machine Learning for IRB Models. EBA/REP/2023/29. EBA.
11. European Parliament. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence. *Official Journal of the European Union*, L 1689/1.
12. Financial Conduct Authority (FCA). (2022). Consumer Duty: Final Rules and Guidance. PS22/9. FCA.
13. Financial Stability Board (FSB). (2022). Supervisory and Regulatory Issues That Merit Authorities' Attention. FSB, Basel.
14. Monetary Authority of Singapore (MAS). (2019). Principles to Promote Fairness, Ethics, Accountability and Transparency (FEAT) in the Use of Artificial Intelligence and Data Analytics in Singapore's Financial Sector. MAS.
15. Office of the Comptroller of the Currency (OCC). (2021). Model Risk Management. *Comptroller's Handbook*. OCC.
16. Philippon, T. (2019). The FinTech opportunity. In T. Beck & R. Levine (Eds.), *Finance, Growth and Inequality*. Edward Elgar Publishing.
17. Prudential Regulation Authority (PRA). (2023). Model Risk Management Principles for Banks. SS1/23. Bank of England.
18. Securities and Exchange Commission (SEC). (2023). Conflicts of Interest Associated with the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers. Proposed Rule, Release No. IA-6353.
19. Wachter, S., Mittelstadt, B. & Russell, C. (2018). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–887.
20. White House. (2023). Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. Executive Order 14110, 30 October 2023.
21. Zetzsche, D.A., Buckley, R.P. & Arner, D.W. (2021). Regulating Libra: The transformative potential of Facebook's cryptocurrency and possible regulatory responses. *Oxford Journal of Legal Studies*, 41(1), 80–113.
22. Zetzsche, D.A., Buckley, R.P., Arner, D.W. & Barberis, J.N. (2017). Regulating a revolution: From regulatory sandboxes to smart regulation. *Fordham Journal of Corporate and Financial Law*, 23(1), 31–103.