International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

Signal & Oversight: Machine Intelligence Meets Financial Regulation

Neha Tyagi

Senior Vice President, Bank of New York, USA

ABSTRACT: The accelerating deployment of machine intelligence across financial markets, credit systems, insurance underwriting, and regulatory surveillance is reshaping the fundamental relationship between financial institutions and their overseers. This paper examines how algorithmic systems — from predictive credit models to high-frequency trading engines and generative AI compliance tools — are transforming both the objects of financial regulation and the instruments through which regulation is conducted.

We identify a structural duality at the heart of contemporary financial governance: machine intelligence simultaneously functions as a source of systemic risk requiring regulatory attention, and as a tool through which supervisory agencies are beginning to exercise oversight more effectively. We term this the Signal-Oversight Duality. Drawing on regulatory economics, information theory, and institutional analysis, we argue that this duality introduces fundamental tensions into existing oversight frameworks that cannot be resolved by incremental adaptation alone.

The paper introduces the concept of Supervisory Machine Intelligence (SMI) — the deployment of algorithmic systems by regulatory bodies for monitoring, enforcement, and market surveillance — and analyses its implications for institutional design, regulatory legitimacy, and financial stability. We propose a tripartite governance framework — Algorithmic Accountability, Supervisory Capacity, and Systemic Resilience — as the basis for a coherent international response to the machine intelligence challenge in finance.

Our analysis draws on verified regulatory documents, peer-reviewed literature, and publicly available supervisory guidance from the Financial Stability Board, the Basel Committee on Banking Supervision, the UK Financial Conduct Authority, the US Securities and Exchange Commission, and the European Securities and Markets Authority.

KEYWORDS: machine learning, financial regulation, supervisory technology, algorithmic trading, SupTech, RegTech, systemic risk, signal extraction, high-frequency trading, financial stability

I. INTRODUCTION

The phrase 'machine intelligence' once belonged to the vocabulary of computer science. In contemporary financial services, it has become an operational reality — embedded in credit decisions, market microstructure, fraud detection, consumer engagement, regulatory reporting, and, increasingly, the supervisory activities of financial regulators themselves. The question is no longer whether machine intelligence will transform financial regulation, but how that transformation should be governed.

Financial markets have always processed signals — price movements, earnings announcements, credit ratings, macroeconomic indicators. What has changed fundamentally is the capacity to extract, process, and act upon signals at speeds and scales that exceed human cognitive reach. A high-frequency trading algorithm may execute thousands of transactions per second; a machine learning credit model may process hundreds of variables simultaneously; a large language model compliance tool may review thousands of regulatory documents in minutes. These capabilities create extraordinary efficiencies and novel risks in equal measure (O'Neil, 2016; Goodfellow, Bengio & Courville, 2016).

Regulators are confronted with a double-edged challenge. On one side, they must supervise institutions whose core operations are increasingly opaque, automated, and fast-moving — outpacing traditional supervisory methodologies built around periodic reporting, annual inspections, and rule-based analysis. On the other side, they are themselves beginning to deploy algorithmic tools — commonly grouped under the label Supervisory Technology, or SupTech — to conduct surveillance, detect anomalies, and enforce compliance at scale (FSB, 2017; BIS, 2018).

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

This paper makes four contributions. First, it characterises the Signal-Oversight Duality — the structural condition in which machine intelligence simultaneously generates the signals that regulators must interpret and constitutes part of the oversight apparatus through which regulators operate. Second, it analyses the distinctive risk categories that machine intelligence introduces into financial systems. Third, it introduces the concept of Supervisory Machine Intelligence and examines its institutional implications. Fourth, it proposes a tripartite governance framework — Algorithmic Accountability, Supervisory Capacity, and Systemic Resilience — to guide the international regulatory response.

The paper proceeds as follows. Section 2 surveys the relevant literature across financial regulation, information economics, and machine learning. Section 3 characterises the principal applications of machine intelligence in financial services. Section 4 introduces the Signal-Oversight Duality and its regulatory implications. Section 5 analyses Supervisory Machine Intelligence. Section 6 examines systemic risk dimensions. Section 7 presents the tripartite governance framework. Section 8 discusses implications. Section 9 concludes.

II. LITERATURE REVIEW

The literature relevant to this paper spans three partially overlapping fields: financial regulation and governance, the economics of information and signalling, and the technical and ethical dimensions of machine learning.

Within financial regulation, the foundational contributions of Stigler (1971) on regulatory capture, and of Fama (1970) on the efficient market hypothesis, establish the twin poles between which contemporary debates are structured. Stigler's insight — that regulatory bodies are susceptible to capture by the industries they oversee — acquires new dimensions when the regulated entities deploy AI systems whose technical complexity creates informational asymmetries that regulators may struggle to overcome. Fama's efficient market framework, meanwhile, is challenged by the demonstration that machine learning systems can, under certain conditions, extract signals from market data that destabilise the informational efficiency the framework presupposes (Goodfellow, Bengio & Courville, 2016).

The economics of information provides essential analytical tools. Akerlof's (1970) seminal analysis of adverse selection in markets with asymmetric information anticipates contemporary concerns about AI-driven credit scoring, in which lenders may deploy models that generate informational advantages over borrowers who cannot interrogate the basis of lending decisions. Spence's (1973) signalling framework similarly illuminates how AI-based credentialing and scoring systems may alter the dynamics of labour market and credit market signalling in ways that carry distributional consequences.

Within the machine learning literature, the foundational textbook of Goodfellow, Bengio and Courville (2016) establishes the technical architecture underpinning contemporary financial AI — particularly deep neural networks, convolutional architectures, and recurrent models — while acknowledging the interpretability challenges these systems present. O'Neil's (2016) widely cited critique of algorithmic decision-making in social contexts, though not limited to financial services, provides an important normative counterpoint, demonstrating how seemingly neutral models may encode and amplify existing social and economic inequalities.

The regulatory response to algorithmic finance has been analysed by Arner, Barberis and Buckley (2017), who identify regulatory technology — RegTech — as a structural response to compliance complexity. Zetzsche et al. (2017) extend this framework, examining how regulatory sandboxes and adaptive regulatory design can accommodate rapid technological change. The FSB's 2017 report on AI and machine learning in financial services provided the first comprehensive international assessment of stability implications, identifying third-party concentration, model opacity, and data governance as primary concerns.

More recently, Wachter, Mittelstadt and Russell (2018) have addressed the legal and ethical dimensions of algorithmic decision-making under data protection frameworks, arguing that post-hoc explanation requirements — as embodied in GDPR Article 22 — may be insufficient to ensure genuine accountability for algorithmic outcomes. Their framework has direct relevance to the financial sector, where credit decisions, insurance pricing, and fraud assessments are increasingly AI-mediated. The FSB's 2024 report on the financial stability implications of AI provides the most current authoritative assessment, identifying generative AI, third-party concentration, and market correlation risks as priorities for supervisory attention.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

The SupTech literature — examining the use of technology by supervisors — is comparatively underdeveloped, reflecting both the relative novelty of the phenomenon and its operational sensitivity. Ellul et al. (2021) provide an early taxonomy of SupTech applications, while the Bank for International Settlements (BIS, 2018) has published technical notes on the use of machine learning in supervisory contexts. This paper builds on this emerging literature by developing a more systematic conceptualisation of Supervisory Machine Intelligence.

III. MACHINE INTELLIGENCE IN FINANCIAL SERVICES: PRINCIPAL APPLICATIONS

3.1 Credit and Lending

Machine learning models have substantially displaced traditional statistical approaches — logistic regression, scorecards — in credit underwriting across both retail and commercial lending. Contemporary credit models exploit alternative data sources including utility payments, rental histories, social media signals, and device usage patterns to extend creditworthiness assessments beyond the populations captured by traditional bureau data (FSB, 2024). The benefits include expanded access to credit for 'thin-file' borrowers and improved default prediction accuracy.

The regulatory challenges are commensurate. Under the Equal Credit Opportunity Act (US), the Fair Housing Act, and equivalent EU non-discrimination frameworks, lenders must be able to demonstrate that credit decisions do not discriminate on protected characteristics. Machine learning models — particularly deep neural network architectures — may produce discriminatory outcomes even when protected characteristics are excluded from the input set, through proxy variables that correlate with protected characteristics (O'Neil, 2016). The Consumer Financial Protection Bureau (CFPB) and the UK FCA have both signalled heightened scrutiny of algorithmic discrimination in lending, with the FCA's Consumer Duty framework (FCA, 2022) imposing positive obligations to demonstrate good outcomes for diverse consumer groups.

3.2 Algorithmic and High-Frequency Trading

Algorithmic trading — and its high-frequency subset — represents the most established application of machine intelligence in financial markets, predating the contemporary ML wave. Algorithmic strategies account for a substantial proportion of equity market volume across major venues, with high-frequency trading firms providing a significant share of liquidity in both equities and derivatives markets (SEC, 2023). The signal-extraction function is central to these strategies: algorithms are trained to detect short-lived informational signals in order flow, price movements, and news sentiment, and to trade on these signals before they are incorporated into prices.

The regulatory concerns are well-documented: market manipulation through layering and spoofing; flash crashes precipitated by correlated algorithmic responses to common signals; and the potential for AI-driven strategies to extract value from retail investors through speed and informational advantages that cannot be replicated. The SEC's Regulation SCI (Systems Compliance and Integrity) and MiFID II's algorithmic trading requirements under Article 17 establish frameworks for operational risk management, but both predate the modern ML era and are widely acknowledged to require updating (SEC, 2023; ESMA, 2021).

3.3 Fraud Detection and AML

Anti-money laundering (AML) and fraud detection represent perhaps the most clearly beneficial and regulatorily endorsed application of machine intelligence in financial services. Machine learning models significantly outperform rule-based transaction monitoring systems in detecting suspicious patterns, reducing both false positive rates — which impose significant compliance costs — and false negative rates — which represent regulatory and reputational risk (FSB, 2017). FinCEN in the United States and the Financial Action Task Force (FATF) internationally have both endorsed the use of advanced analytics, including machine learning, in AML/CFT compliance programmes.

The regulatory framework governing AML/CFT — rooted in the Bank Secrecy Act (US), the Fourth and Fifth Anti-Money Laundering Directives (EU), and FATF Recommendations — was not designed with AI-mediated compliance in mind. Questions of model explainability, auditability, and human oversight arise acutely when Suspicious Activity Reports are generated by algorithmic systems: can a compliance officer meaningfully certify a SAR they did not independently generate? Can a regulator audit an AML system they cannot interrogate?

3.4 Insurance Underwriting and Pricing

The insurance sector has adopted machine learning for risk pricing, claims processing, fraud detection, and customer lifecycle management. Granular risk pricing — enabled by telematics data, wearable devices, and behavioural data

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

streams — challenges the pooling principles that underpin traditional insurance models, with significant distributional implications (O'Neil, 2016). Regulatory frameworks across major jurisdictions impose constraints on the variables that may be used in insurance pricing — prohibiting gender in the EU following the Test-Achats judgment — but the opacity of ML models creates enforcement challenges that conventional actuarial oversight mechanisms are ill-equipped to address.

3.5 Compliance and Regulatory Reporting

Most recently, financial institutions have begun deploying large language models and other generative AI tools within compliance functions: for regulatory change management, contract review, policy gap analysis, and regulatory report drafting. This deployment introduces a recursive dimension to the machine intelligence challenge: the tools being used to ensure compliance with AI regulation are themselves AI systems, subject to the same concerns about reliability, hallucination, auditability, and accountability that pervade their application elsewhere (FSB, 2024).

IV. THE SIGNAL-OVERSIGHT DUALITY

We introduce the Signal-Oversight Duality as a conceptual framework for understanding the structural position of machine intelligence in contemporary financial regulation. The duality has two dimensions.

The Signal Dimension captures the role of machine intelligence as a signal-processing apparatus embedded in financial institutions. Financial markets are, fundamentally, signal-processing systems: prices aggregate dispersed information; credit scores condense complex risk profiles; trading algorithms detect and exploit informational inefficiencies. Machine intelligence dramatically amplifies the signal-processing capacity of financial institutions, extracting value from data volumes and at speeds that are beyond human analytical reach. From a regulatory perspective, the signals processed by these systems are both the objects of oversight — the decisions, transactions, and risk assessments that regulators must supervise — and the medium through which oversight is exercised, in the form of supervisory data and reporting.

The Oversight Dimension captures the role of machine intelligence as a supervisory apparatus deployed by regulatory bodies. As financial institutions' operations have grown in complexity and velocity, regulators have recognised that traditional oversight methodologies — periodic reporting, annual examinations, retrospective enforcement — are insufficient. SupTech — supervisory technology — encompasses the deployment of machine learning, natural language processing, network analysis, and other AI-adjacent tools by supervisors to conduct real-time surveillance, anomaly detection, and enforcement analysis (BIS, 2018).

The duality creates a structural interdependency: regulators increasingly depend on algorithmic systems to supervise algorithmic systems. This creates at least three forms of institutional tension. First, a competence asymmetry: financial institutions deploying frontier AI systems will typically possess greater technical sophistication than the supervisors overseeing them, creating informational advantages that may impede effective oversight. Second, a legitimacy tension: when regulatory enforcement actions are triggered or supported by algorithmic analysis, the procedural fairness norms that underpin regulatory legitimacy — the right to be heard, to understand the basis of decisions, and to contest them — may be inadequately satisfied. Third, a systemic feedback risk: if both regulated institutions and their supervisors deploy similar or correlated algorithmic systems, the potential for self-reinforcing feedback loops — in which supervisory interventions and institutional responses are both algorithmically mediated — creates novel forms of systemic fragility.

V. SUPERVISORY MACHINE INTELLIGENCE

5.1 Definition and Scope

We define Supervisory Machine Intelligence (SMI) as the deployment of algorithmic systems — including but not limited to machine learning models, natural language processing tools, network analysis platforms, and generative AI applications — by financial regulatory and supervisory authorities, for the purposes of market surveillance, entity supervision, consumer protection, enforcement, and systemic risk monitoring.

SMI encompasses a range of existing and emerging SupTech applications, including: automated transaction monitoring systems used by financial intelligence units; market surveillance platforms deployed by securities regulators; regulatory reporting analysis tools used by prudential supervisors; and AI-assisted examination preparation systems. The scale of

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

SupTech investment has accelerated significantly since the FSB's 2017 assessment: the BIS reports that over 60 jurisdictions have active SupTech programmes as of 2024 (BIS, 2018; FSB, 2024).

5.2 Benefits of SMI

The case for SMI rests on three pillars. First, scalability: the volume of financial data now generated by regulated institutions — transaction records, communications, risk reports, market data — vastly exceeds what human supervisory teams can review manually. SMI enables regulators to monitor a far greater proportion of financial activity than was previously feasible. Second, velocity: algorithmic surveillance can identify suspicious patterns or threshold breaches in near real-time, enabling preventive or early intervention rather than retrospective enforcement. Third, consistency: algorithmic rules, applied uniformly across institutions, reduce the scope for differential treatment and supervisory arbitrariness that may arise in manual examination processes.

5.3 Institutional Design Challenges

Despite these benefits, SMI raises profound institutional design challenges that have received insufficient attention in the regulatory literature. We identify four primary challenges.

The first is the explainability-accountability nexus. When an SMI system generates an alert, triggers an investigation, or contributes to an enforcement recommendation, the procedural fairness obligations owed to the regulated entity require that the basis for supervisory action can be meaningfully articulated and contested. Algorithmic outputs that cannot be explained in human-intelligible terms may be incompatible with constitutional due process norms and administrative law principles that require reasoned decision-making (Wachter, Mittelstadt & Russell, 2018).

The second is the audit trail requirement. Regulatory enforcement actions must be capable of withstanding legal challenge. An SMI system that produces evidential outputs — flagging a market participant as having engaged in manipulation, for example — must generate an audit trail that is both technically rigorous and legally admissible. The standards for algorithmic evidence in enforcement proceedings remain underdeveloped across major jurisdictions.

The third is the feedback and gaming risk. If market participants can infer the parameters and decision rules of SMI surveillance systems, they may structure their activities to avoid detection — a form of regulatory arbitrage enabled by algorithmic transparency. This creates a tension between the transparency obligations that democratic accountability imposes on public regulators and the operational security that effective surveillance requires.

The fourth is concentration and vendor risk. SMI systems are frequently procured from a small number of specialist vendors. This creates systemic dependencies analogous to those identified in the financial institution context: if a supervisory technology vendor experiences a failure, the oversight capacity of multiple regulators across multiple jurisdictions may be simultaneously degraded (FSB, 2024). The procurement of supervisory technology from the same vendor base as regulated institutions' compliance technology creates further interdependencies.

VI. SYSTEMIC RISK DIMENSIONS OF MACHINE INTELLIGENCE IN FINANCE

The FSB's 2024 report identifies four AI-related vulnerabilities with systemic risk potential: third-party concentration, market correlations, cyber risk, and model risk. We extend this framework by identifying two additional dimensions: supervisory capacity erosion and the legitimacy deficit.

6.1 Third-Party Concentration

A small number of cloud providers and AI infrastructure vendors — predominantly large US technology firms — supply a disproportionate share of the computational infrastructure and foundational model capabilities on which both financial institutions and their supervisors now depend. The FSB has characterised this concentration as a macro-level systemic risk: an operational failure, security incident, or commercial disruption affecting a major provider could simultaneously impair risk management, trading, compliance, and supervisory functions across large numbers of institutions and jurisdictions (FSB, 2024). Existing regulatory frameworks — including DORA in the EU — address ICT concentration risk at the entity level, but lack adequate macro-prudential tools for managing systemic concentration across the financial system as a whole.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

6.2 Market Correlation and Herding

Machine learning models trained on similar datasets and optimising similar objectives may generate correlated outputs across institutions. In credit markets, models trained on similar historical default data may simultaneously tighten lending standards during periods of economic stress, amplifying credit contraction. In trading, models trained on similar technical signals may generate correlated position-taking or risk-reduction behaviour, amplifying price moves and reducing the market diversity that supports liquidity (Goodfellow, Bengio & Courville, 2016; FSB, 2017). This herding dynamic creates a form of endogenous systemic risk that does not arise from individual institutional failures but from the aggregate behaviour of the algorithmic ecosystem.

6.3 Supervisory Capacity Erosion

As financial institutions' operations become more complex and algorithmically mediated, the gap between the technical sophistication required to supervise those operations and the capabilities of regulatory staff widens. This supervisory capacity erosion represents a systemic risk not of the conventional kind — it does not directly threaten the solvency of any institution — but rather a meta-risk: the erosion of the oversight capacity on which the entire regulatory framework depends. If supervisors cannot interrogate the AI systems they are charged with overseeing, enforcement becomes nominal rather than substantive, and the risk management frameworks mandated by regulators may be implemented without meaningful external scrutiny.

6.4 The Legitimacy Deficit

The deployment of AI in consequential financial decisions — credit refusals, insurance exclusions, fraud flags, enforcement referrals — creates legitimacy challenges that extend beyond technical risk management. When individuals are adversely affected by algorithmic decisions they cannot understand or contest, the social legitimacy of the financial system and its regulatory framework is undermined. O'Neil (2016) characterises this as a structural feature of large-scale algorithmic governance, not an incidental failure of individual systems. The legitimacy deficit has regulatory significance: systems that are perceived as opaque and arbitrary will generate political pressure for blunt interventions — such as algorithmic prohibitions or excessive disclosure requirements — that may destroy value without addressing the underlying problem.

VII. A TRIPARTITE GOVERNANCE FRAMEWORK

In response to the challenges identified above, we propose a tripartite governance framework comprising three interlocking pillars: Algorithmic Accountability, Supervisory Capacity, and Systemic Resilience. The framework is intended to operate at both the institutional level — governing the behaviour of individual financial firms and their AI systems — and the systemic level — governing the macro-prudential and international dimensions of machine intelligence in finance.

7.1 Pillar I — Algorithmic Accountability

Algorithmic Accountability encompasses the obligations of financial institutions to ensure that AI systems deployed in consequential financial decisions are transparent, explainable, auditable, and contestable. The pillar builds on existing frameworks — GDPR Article 22, the EU AI Act's high-risk system requirements, the UK FCA's Consumer Duty, and the OCC's model risk management guidance — but extends them in two important directions.

First, it requires model cards and algorithmic impact assessments as standard components of pre-deployment governance for AI systems used in credit, insurance, trading, and compliance contexts. These documents — analogous to the environmental impact assessments used in project finance — should specify the training data, performance metrics, bias testing results, explainability approach, and ongoing monitoring arrangements for each deployed system.

Second, it requires institutions to maintain human decision-override capacity for all AI-mediated decisions affecting individual consumers or counterparties. The principle of meaningful human oversight — enshrined in the EU AI Act for high-risk applications — must be operationalised as a genuine governance requirement, not a nominal check-box exercise. This requires investment in the expertise of human reviewers and the design of decision interfaces that provide genuine decision support rather than cognitive rubber-stamping.

7.2 Pillar II — Supervisory Capacity

Supervisory Capacity encompasses the institutional investment required to ensure that regulatory bodies can effectively oversee AI-driven financial systems. This pillar has three components.

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

The first is technical expertise: regulators must recruit, develop, and retain staff with meaningful AI and data science competencies. The FSB, BIS, and IOSCO should support the development of a shared international supervisory talent framework, enabling smaller jurisdictions to access AI expertise they could not independently resource.

The second is SupTech investment: regulators must invest in SMI capabilities that enable real-time, data-driven supervision across the full range of regulated activity. However, this investment must be accompanied by robust governance frameworks for SMI itself — ensuring that supervisory algorithms are themselves explainable, auditable, and subject to independent oversight.

The third is regulatory sandboxes for supervision: regulators should pilot new supervisory approaches — including AI-based examination tools, automated reporting analysis systems, and algorithmic market surveillance platforms — in controlled environments before deploying them at scale, following a 'test and learn' approach analogous to the innovation sandboxes they provide for regulated entities.

7.3 Pillar III — Systemic Resilience

Systemic Resilience encompasses the macro-prudential measures required to manage the systemic risks identified in Section 6. It has four components.

Concentration oversight: regulators should extend macro-prudential surveillance to the third-party AI and cloud infrastructure market, treating excessive concentration in supervisory technology providers as a systemic risk requiring regulatory attention. This may require extending the scope of systemic importance frameworks — currently focused on financial institutions — to cover critical technology providers.

Diversity requirements: regulators should consider whether diversity requirements — mandating that systemically important financial institutions do not rely exclusively on a single AI vendor or model architecture for critical functions — are warranted, drawing on the analogy of diversity requirements in derivatives clearing and settlement.

International coordination: the systemic risks of machine intelligence in finance are inherently cross-border. The FSB, BCBS, and IOSCO should develop international standards for AI governance in financial services that provide a basis for mutual recognition, reducing the compliance costs of jurisdictional fragmentation while establishing common minimum standards for systemic risk management.

Continuous stress testing: financial institutions' AI systems should be subject to periodic adversarial stress tests — examining their behaviour under conditions of market stress, data shift, and adversarial attack — alongside the capital stress tests already conducted by prudential supervisors. The BCBS's model risk framework should be updated to incorporate AI-specific stress scenarios.

VIII. IMPLICATIONS FOR PRACTITIONERS AND POLICYMAKERS

8.1 For Financial Institutions

The tripartite framework has immediate implications for senior leadership in financial institutions. Chief Risk Officers and Chief Compliance Officers should initiate comprehensive AI governance audits — assessing whether deployed AI systems meet the explainability, auditability, and human oversight standards that the framework prescribes. Boards should receive regular reporting on AI-related risk exposures, including model risk, third-party concentration risk, and regulatory compliance risk across all relevant jurisdictions.

Institutions should also invest in the internal capability required to engage meaningfully with regulators on AI governance. The technical complexity of modern AI systems means that regulatory dialogue — whether in examination contexts, consultation responses, or enforcement proceedings — increasingly requires specialist expertise within the compliance and legal functions, not merely technical teams. The development of 'regulatory AI literacy' across the second and third lines of defence is an institutional priority that deserves board-level attention.

8.2 For Regulators and International Bodies

For regulatory bodies, the analysis reinforces the urgency of the SupTech investment agenda. Supervisors who cannot interrogate the systems they are overseeing are supervisors in name only. The FSB's recommendation to close data gaps

International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

and enhance supervisory capabilities, set out in its 2024 report, should be treated not as a medium-term aspiration but as an immediate operational priority (FSB, 2024).

International bodies — the FSB, BCBS, IOSCO, and the IMF — should accelerate work on the cross-border dimensions of AI governance in finance. The jurisdictional fragmentation of AI regulation, documented in the previous paper in this series, imposes compliance costs on internationally active firms and creates opportunities for regulatory arbitrage. A minimum international standard for AI governance in systemically important financial institutions would provide a valuable floor, compatible with jurisdictional diversity in implementation above that floor.

IX. CONCLUSION

Machine intelligence has become structural to financial services — not a peripheral innovation but an integral component of how financial institutions make decisions, manage risk, and interact with markets and customers. The regulatory response to this transformation has been substantial but remains incomplete. Frameworks designed for an era of rule-based systems and human decision-makers are being stretched to accommodate AI systems that are faster, more complex, and more opaque than anything their architects envisaged.

This paper has argued that the fundamental challenge is captured by the Signal-Oversight Duality: machine intelligence is simultaneously the phenomenon that financial regulators must oversee and the instrument through which that oversight must increasingly be conducted. This duality creates institutional tensions — around competence, legitimacy, and systemic feedback — that cannot be resolved through technical fixes alone. They require deliberate institutional design.

The tripartite governance framework proposed here — Algorithmic Accountability, Supervisory Capacity, and Systemic Resilience — provides a structured basis for that design. Its three pillars are mutually reinforcing: accountability without supervisory capacity produces unenforceable obligations; supervisory capacity without systemic resilience addresses institutional risks while leaving macro-prudential vulnerabilities unaddressed; resilience without accountability leaves the distributional and legitimacy dimensions of machine intelligence in finance unresolved.

The signal is complex and the horizon uncertain. But the imperative to develop governance frameworks adequate to the challenge of machine intelligence in finance is clear, immediate, and of profound importance to the stability and fairness of the global financial system. The question is not whether to regulate machine intelligence in finance — the consequences of failing to do so are too serious — but how to regulate it with sufficient rigour, intelligence, and agility to keep pace with the systems being governed.

REFERENCES

1. Akerlof, G.A. (1970). The market for 'lemons': Quality uncertainty and the market mechanism. Quarterly Journal of Economics, 84(3), 488–500. https://doi.org/10.2307/1879431
2. 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.
3. Bank for International Settlements (BIS). (2018). Sound Practices: Implications of Fintech Developments for Banks and Bank Supervisors. Basel Committee on Banking Supervision. BIS.
4. Consumer Financial Protection Bureau (CFPB). (2022). Circular 2022-03: Adverse Action Notification Requirements in Connection with Credit Decisions Based on Complex Algorithms. CFPB.
5. Ellul, A., Jotikasthira, C., Lundblad, C. & Wang, Y. (2021). Review of SupTech and RegTech developments. Journal of Financial Regulation, 7(1), 1–30.
6. European Securities and Markets Authority (ESMA). (2021). Final Report on Algorithms and High-Frequency Trading. ESMA70-156-5533. ESMA.
7. Fama, E.F. (1970). Efficient capital markets: A review of theory and empirical work. Journal of Finance, 25(2), 383–417. https://doi.org/10.2307/2325486
8. Financial Conduct Authority (FCA). (2022). Consumer Duty: Final Rules and Guidance. PS22/9. FCA.
9. Financial Stability Board (FSB). (2017). Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications. FSB, Basel.
10. Financial Stability Board (FSB). (2024). The Financial Stability Implications of Artificial Intelligence. FSB, Basel. November 2024.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 4, July-August 2025||

DOI:10.15662/IJRPETM.2025.0804015

11. Goodfellow, I., Bengio, Y. & Courville, A. (2016). Deep Learning. MIT Press.
12. O'Neil, C. (2016). Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Crown Publishers.
13. Prudential Regulation Authority (PRA). (2023). Model Risk Management Principles for Banks. SS1/23. Bank of England.
14. 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.
15. Spence, M. (1973). Job market signalling. Quarterly Journal of Economics, 87(3), 355–374. https://doi.org/10.2307/1882010
16. Stigler, G.J. (1971). The theory of economic regulation. Bell Journal of Economics and Management Science, 2(1), 3–21. https://doi.org/10.2307/3003160
17. 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.
18. White House. (2023). Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. Executive Order 14110, 30 October 2023.
19. 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.