

Artificial Intelligence and Financial-Sector Supervision in Africa: Opportunities, Risks and a Framework for Responsible Adoption.

Akomolehin Francis Olugbenga

<https://orcid.org/0009-0000-4896-9230>

Department of Finance, College of Management & Social Science Afe Babalola University, Ado - Ekiti
Ekiti

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.150700084>

Received: 29 July 2026; Accepted: 03 August 2026; Published: 13 August 2026

ABSTRACT

Artificial intelligence (AI) is increasingly used in financial services and supervision, yet financial institutions and FinTech firms are adopting it faster than many African regulators can effectively oversee it. This study examined the applications, opportunities, risks and regulatory conditions associated with AI-enabled financial-sector supervision in Africa. It combined a systematic literature review with comparative regulatory analysis covering publications issued between January 2020 and July 2026. Evidence was retrieved from Scopus, Web of Science, SSRN, the IMF, World Bank, Bank for International Settlements, African Development Bank and relevant African regulatory authorities. Following PRISMA procedures, 565 records were identified, 141 duplicates were removed, and 424 records were screened. After full-text and quality assessments, 56 publications were included in the final thematic synthesis. Regulatory arrangements in Nigeria, Ghana, Kenya and South Africa were compared across AI policy, SupTech initiatives, data protection, cybersecurity, consumer protection, regulatory sandboxes, institutional capacity and accountability. The findings showed that AI could strengthen fraud detection, risk-based supervision, early-warning systems, regulatory-reporting analysis and consumer monitoring. However, poor data quality, algorithmic bias, limited explainability, privacy and cybersecurity risks, skills shortages and vendor dependence constrained responsible adoption. South Africa demonstrated the strongest disclosed financial-sector AI readiness, while Kenya had the clearest national AI strategy. Nigeria and Ghana had important digital-finance and cybersecurity foundations but limited publicly documented AI-enabled supervisory applications. The study developed an African Responsible AI–Financial Supervision Framework comprising six connected pillars and a phased adoption cycle. It recommends proportionate implementation, common model-validation standards, stronger inter-agency cooperation, shared regional infrastructure and independent oversight.

Keywords: artificial intelligence; financial-sector supervision; SupTech; responsible AI; financial regulation; Africa

INTRODUCTION

1.1 Background to the Study

Artificial intelligence (AI) has moved rapidly into mainstream financial services. Banks, insurers, securities firms and FinTech companies now apply machine learning, natural-language processing and predictive analytics to credit assessment, fraud detection, customer service, trading, compliance and risk management. A recent review found that forecasting, classification, early-warning systems and large-scale data analysis are among the dominant AI applications in finance (Bahoo et al., 2024). This expansion has changed the volume, speed and complexity of risks confronting financial authorities.

Financial supervision has therefore begun to shift from periodic returns, manual reviews and retrospective inspections towards continuous, data-driven oversight. Supervisory technology, commonly known as SupTech,

refers to advanced technologies used by regulatory authorities to support data collection, analysis and supervisory decision-making. SupTech systems can process regulatory returns, analyse complaints and transactions, identify abnormal patterns, assess institutional risk and direct supervisory attention to areas requiring early intervention. A review of 71 SupTech tools across 20 jurisdictions found that these applications improved the analysis of quantitative and qualitative information, although their effectiveness depended on data quality, staff competence and their integration with supervisory judgement (Beerman et al., 2021).

AI extends these capabilities. It can detect suspicious transactions, classify institutions according to risk, identify inconsistencies in regulatory reports, monitor market conduct and generate early-warning indicators. It may also support consumer protection by analysing complaints, advertisements, product terms and digital transaction patterns. These applications can help regulators move from compliance-driven supervision towards more preventive and risk-based approaches, especially where supervisory responsibilities are expanding faster than available staff and budgets (Dohotaru et al., 2025).

Interest in SupTech has consequently increased among central banks and other financial authorities. Yet experimentation has advanced faster than full institutional use. Prenio (2024) found that many SupTech tools remained insufficiently embedded in routine supervision because authorities lacked clear strategies, reliable data, appropriate skills and sustained engagement between supervisory and technology teams. Similarly, a survey of 55 financial authorities identified poor data quality, limited staff resources and increasingly complex cyber, climate and technology risks as common supervisory difficulties (Crisanto et al., 2024).

African regulators occupy an important but challenging position in this development. Digital banking, mobile money and FinTech activities have expanded rapidly across the continent, increasing the need for timely and technology-supported oversight. AI research in Africa has also grown, although it remains unevenly distributed and is led largely by South African institutions (Kondo & Diwani, 2023). More directly, a World Bank survey of 27 emerging-market financial authorities, including 17 African authorities, found that financial institutions and FinTech firms generally remained ahead of regulators in AI adoption. Almost half of the surveyed authorities reported that they lacked sufficient systems, resources and processes to monitor AI developments and their effects (World Bank, 2025).

1.2 Research Problem

The central problem is the widening gap between AI adoption in financial markets and the capacity of regulators to understand, supervise and use the technology responsibly. Financial institutions increasingly apply AI to decisions affecting credit, pricing, fraud controls and customer access, while many African regulators still depend on fragmented reporting systems, delayed data and labour-intensive supervision. Skills shortages, weak interoperability, cybersecurity exposure and constrained technology budgets further limit regulatory responses.

AI-assisted supervision could improve regulatory efficiency, but it also creates difficult governance questions. Models trained on incomplete or unrepresentative data may reproduce discrimination. Complex algorithms may produce decisions that supervisors, regulated firms and affected consumers cannot adequately explain. Expanded access to granular financial data raises privacy and cybersecurity risks, while dependence on external technology vendors can weaken regulatory control. Accountability also becomes unclear when an automated alert, risk classification or enforcement decision proves inaccurate. Peer-reviewed evidence confirms that privacy, algorithmic bias, transparency, fairness and regulatory oversight remain central concerns in AI-enabled banking (Fundira & Mbohwa, 2025). Existing knowledge offers limited guidance for resolving these issues in Africa. Much of the literature focuses on commercial uses of AI or regulatory developments in advanced economies. African studies tend to address financial inclusion, digital lending or general AI adoption, with little direct attention to how financial authorities use AI or supervise its use. Comparative evidence on Nigeria, Ghana, Kenya and South Africa is particularly limited, despite differences in their financial markets, data-protection laws, cybersecurity arrangements, regulatory sandboxes and institutional capacities.

1.3 Research Gap and Motivation

The study addressed four connected gaps. First, the **contextual gap** arose from the limited attention given to African central banks and financial regulators in AI-supervision research. Second, the **comparative gap** concerned the absence of structured comparisons of national approaches within Africa. Third, the **conceptual gap** reflected the tendency to discuss AI's efficiency benefits separately from its consumer, ethical and institutional risks. Recent scholarship has similarly observed that regulatory, ethical and consumer dimensions are often treated in isolation (Fundira & Mbohwa, 2025). Fourth, the **policy gap** concerned the lack of an operational framework that translates responsible-AI principles into steps suited to African regulatory conditions. These gaps motivated a systematic review of recent evidence and a comparative analysis of Nigeria, Ghana, Kenya and South Africa. The study integrated AI applications, supervisory benefits, technological risks and institutional readiness, using the findings to develop an African Responsible AI–Financial Supervision Framework.

1.4 Research Questions

The study answered the following questions:

1. How is AI being applied or proposed for financial-sector supervision in Africa?
2. What opportunities does AI offer for fraud detection, risk-based supervision and consumer protection?
3. What risks and constraints affect its responsible use by African financial authorities?
4. How do the regulatory approaches of Nigeria, Ghana, Kenya and South Africa compare?
5. What principles and institutional arrangements should guide responsible AI adoption in African financial-sector supervision?

1.5 Research Objectives

The study examined current and emerging applications of AI in African financial-sector supervision; assessed its opportunities for supervisory effectiveness; identified associated risks and institutional constraints; compared relevant regulatory approaches in Nigeria, Ghana, Kenya and South Africa; and developed an African Responsible AI–Financial Supervision Framework.

1.6 Significance of the Study

The study provides central banks and financial regulators with a compact basis for assessing where AI can strengthen supervision and what safeguards should accompany its use. It also supports deposit insurers, securities regulators and data-protection authorities in coordinating responses to model risk, consumer harm, privacy and cybersecurity. Financial institutions and FinTech firms may use the framework to anticipate supervisory expectations concerning data governance, explainability and accountability. For the African Union, the African Development Bank and other development institutions, the country comparison identifies areas suitable for technical assistance, shared supervisory infrastructure and regional policy coordination. Finally, the study extends research on AI governance and financial stability by combining recent peer-reviewed evidence with African regulatory experience and by treating responsible adoption as a practical institutional issue rather than a set of abstract ethical principles.

LITERATURE REVIEW

2.1 Conceptual Review

2.1.1 Artificial Intelligence in Financial Services

Artificial intelligence refers to machine-based systems that infer from supplied data how to generate predictions, recommendations, content or decisions with varying levels of autonomy and adaptability

(Organisation for Economic Co-operation and Development [OECD], 2024). In financial services, AI describes a range of related technologies rather than one application.

Machine learning allows systems to identify patterns and improve predictions from historical data. Banks and FinTech firms use it for credit scoring, transaction classification, default prediction and fraud detection. Natural-language processing converts unstructured materials, such as financial reports, customer complaints and regulatory filings, into information that can be searched and analysed. Predictive analytics estimates future outcomes, including credit deterioration, liquidity stress or institutional failure, while anomaly-detection systems identify transactions and behaviour that depart from established patterns. Bahoo et al. (2024) found that forecasting, classification, detection, early-warning systems and large-scale data analysis were among the main applications of AI in finance. These capabilities have increased operational efficiency, but they have also made financial activities more complex and data-dependent.

2.1.2 Financial-Sector Supervision

Financial-sector supervision comprises the activities through which public authorities monitor financial institutions, markets and products to preserve soundness, fair conduct and public confidence. **Prudential supervision** assesses capital adequacy, liquidity, governance, asset quality and risk-management practices to reduce the likelihood and cost of institutional failure. **Conduct supervision** examines how financial firms design, market and deliver products, including whether they treat customers fairly and disclose material information.

Market surveillance monitors trading, pricing and transaction patterns to identify manipulation, insider dealing, money laundering and other forms of market abuse. **Consumer-protection supervision** addresses misleading disclosures, unfair charges, discriminatory decisions, aggressive digital lending and misuse of customer data. **Financial-stability oversight** considers risks that extend beyond an individual institution, including interconnected exposures, common dependence on technology providers and market-wide reactions to similar AI-generated signals. The boundaries between these functions have become less distinct as digital platforms, mobile finance and automated decisions create risks that may simultaneously affect institutions, consumers and the wider financial system.

2.1.3 AI-Enabled Financial Supervision and SupTech

Supervisory technology, or SupTech, refers to technology used by financial authorities to support supervisory work. AI-enabled SupTech represents its more advanced analytical dimension. It moves supervision beyond the digital collection of returns towards automated interpretation, prediction and prioritisation.

Regulators can use natural-language processing to examine regulatory returns, audit reports, board minutes and consumer complaints. Machine-learning models can compare reported figures across periods and institutions, identify inconsistent submissions and flag unusual changes. Anomaly detection can expose suspicious payments, possible market abuse or transaction networks associated with financial crime.

Predictive systems can combine capital, liquidity, profitability and governance indicators to identify emerging institutional risks and help determine which firms require inspection. Beerman et al. (2021) found that SupTech tools supported both quantitative and qualitative risk assessments, although their outputs needed to complement rather than replace supervisory judgement. Prenio (2024) similarly showed that technical deployment alone did not ensure effective supervision; the tools had to be embedded in supervisory processes and regularly used by responsible staff.

2.1.4 Opportunities Presented by AI

AI offers six closely related opportunities for African financial authorities. First, in **fraud and financial-crime detection**, algorithms can analyse large transaction volumes, identify abnormal behaviour and map

connections that manual reviews may miss. This permits quicker investigation of fraud, money laundering and suspicious fund movements.

Second, AI can strengthen **risk-based supervision** by classifying institutions according to their risk profiles and directing attention towards those with higher or rapidly changing exposure. This is particularly useful where the number and complexity of regulated entities exceed available supervisory resources. Dohotaru et al. (2025) observed that AI could assist supervisors in processing unstructured and high-volume information across consumer protection, credit risk, stress testing and anti-money-laundering activities.

Third, **real-time monitoring and early-warning systems** can identify deteriorating liquidity, abnormal withdrawals, rising complaints or unusual market behaviour before conventional reporting cycles reveal them. Fourth, automated **regulatory-reporting analysis** can improve validation, reconciliation and comparison of returns. Fifth, **consumer-protection monitoring** can analyse complaints, advertisements, digital-loan terms and transaction patterns to detect unfair treatment, hidden charges and emerging scams. Finally, automation can release supervisors from repetitive screening tasks and allow scarce professional staff to focus on investigation and judgement. These gains are especially relevant in resource-constrained jurisdictions, although they depend on sound data and operational integration (World Bank, 2025).

2.1.5 Risks Associated with AI-Assisted Supervision

The use of AI does not remove supervisory risk; it changes its form. **Algorithmic bias** may occur when training data underrepresent certain regions, genders, income groups or informal businesses. A model can therefore appear technically accurate while systematically disadvantaging poorly represented groups. **Poor data quality** can also produce false alerts, inaccurate institutional rankings and misleading predictions. These risks are serious in countries where reporting remains fragmented or inconsistent.

Limited explainability makes it difficult for supervisors to understand why a complex model generated a particular risk score or recommendation. This weakens validation and may prevent affected institutions or consumers from challenging a decision. **Data-privacy risks** arise because AI systems often require large volumes of personal and transactional information. Centralising such information can also increase **cybersecurity exposure**, including unauthorised access, data manipulation and adversarial attacks on supervisory models.

Model risk includes faulty assumptions, overfitting, performance deterioration and inappropriate use outside the conditions for which a model was developed. Dependence on cloud providers, external models and specialist consultants creates **vendor risk**, particularly where regulators cannot inspect proprietary algorithms or easily replace a provider. The Financial Stability Board (2024) warned that AI could amplify third-party concentration, cyber risk, model risk and correlated market behaviour.

Automated systems also complicate **accountability**. Responsibility must remain identifiable when an AI recommendation contributes to delayed intervention, wrongful investigation or consumer harm. A further African concern is **exclusion through invisibility**: informal providers, cash-based businesses and consumers with limited digital records may fall outside data-driven monitoring. As a result, AI may produce a detailed picture of the formal digital sector while leaving important sources of vulnerability unobserved.

2.1.6 Responsible AI Adoption

Responsible AI adoption means that supervisory authorities pursue technological efficiency within enforceable legal, ethical and operational limits. **Legality** requires AI use to comply with supervisory mandates, data-protection laws and procedural rights. **Fairness** requires testing for unjustified differences across institutions and consumer groups. **Transparency** requires authorities to disclose relevant purposes, data sources and governance arrangements, while **explainability** requires supervisory staff to understand and communicate the main reasons behind consequential outputs.

Proportionality means that controls and levels of automation should reflect the risk and consequence of each application. A tool used to organise documents may require lighter controls than one used to classify a bank as high risk. **Security** requires protection of data, models and infrastructure throughout the AI lifecycle. **Accountability** assigns clear responsibility to officials and governing bodies, even where an external vendor supplies the system. **Contestability** provides a process for regulated entities or affected persons to question, correct and seek review of an AI-supported decision.

These principles should operate as governance practices rather than broad declarations. Papagiannidis et al. (2025) organised responsible AI governance around structural, relational and procedural practices, while Fundira and Mbohwa (2025) stressed the need to integrate fairness, privacy, transparency and regulatory oversight. The updated OECD AI Principles similarly place human rights, robustness, security and accountability at the centre of trustworthy AI (OECD, 2024).

2.1.7 Institutional Capacity

Institutional capacity determines whether responsible adoption can move beyond policy statements. Regulators require staff who understand finance, data science, model validation, consumer protection and cybersecurity. Technical expertise must work alongside supervisory knowledge because a statistically sophisticated tool may still misinterpret institutional behaviour or regulatory requirements.

Adequate infrastructure is also necessary for secure data storage, system interoperability, processing capacity and operational continuity. Strong data governance should define data ownership, quality standards, access rights, retention rules and responsibility for corrections. Funding must cover maintenance, retraining, cybersecurity testing and independent validation, not merely initial acquisition.

Inter-agency cooperation is equally important. Central banks, securities and insurance regulators, deposit insurers, financial-intelligence units, competition authorities and data-protection agencies may hold different parts of the information required for effective oversight. Secure information-sharing arrangements can reduce fragmentation while preserving lawful access controls. Finally, regulators need timely and reliable supervisory data. A BIS survey found that limited staff resources, poor data quality and ineffective use of existing information remained common barriers to SupTech development (Crisanto et al., 2024). Institutional readiness therefore conditions the relationship between AI adoption and supervisory effectiveness: where skills, governance and data are weak, greater automation may accelerate existing weaknesses rather than correct them.




Figure 1. Conceptual model of responsible AI-enabled financial-sector supervision.

Source: Authors' conceptualisation based on the reviewed literature.

2.2 Theoretical Review

The study was anchored in the Technology–Organisation–Environment (TOE) Framework and Institutional Theory. Together, these perspectives explain both the practical conditions affecting AI adoption and the external pressures shaping the decisions of financial regulators.

The TOE framework explains organisational technology adoption through three connected contexts: technology, organisation and environment. Although originally developed for business organisations, it applies to central banks and other financial authorities because they make institutional decisions about acquiring, testing and integrating new technologies. Recent AI-adoption studies have retained the framework because it captures organisational and environmental conditions that individual-level acceptance theories do not adequately address (Khanfar et al., 2024).

The technological context concerns the availability and suitability of resources required for AI-assisted supervision. Relevant conditions include data availability and quality, secure processing capacity, interoperability, cybersecurity, system compatibility and model reliability. AI cannot produce credible supervisory assessments when regulatory returns are incomplete, reporting formats are inconsistent or institutional databases cannot communicate. New applications must also connect with existing regulatory-reporting platforms and supervisory procedures. Their continued reliability depends on model validation, performance monitoring and protection against technical failure or manipulation. The perceived usefulness of AI is therefore insufficient to guarantee adoption without a dependable data and technology foundation.

The organisational context covers a regulator's internal ability to adopt and sustain AI. It includes leadership support, staff competence, funding, governance arrangements and institutional readiness. AI-assisted supervision requires teams that combine knowledge of regulation with data science, information technology, cybersecurity and model-risk management. Leadership must define acceptable uses, assign responsibility and support cooperation between supervisory and technical departments. Evidence from public-sector AI adoption shows that managerial attention, specialist skills, data governance and alignment between technology and institutional priorities influence whether AI moves beyond isolated experimentation (Alshahrani et al., 2022). A successful pilot may therefore fail to become part of routine supervision when staff do not trust its outputs or when responsibility for maintenance remains uncertain.

The environmental context comprises factors outside the regulatory institution. These include legislation, data-protection requirements, political conditions, financial-market complexity, national digital infrastructure and international standards. A regulator supervising banks, mobile-money operators, FinTech firms and digital lenders faces different technology and data demands from an authority responsible for a less diversified market. External technical assistance and regulatory cooperation may support adoption, while unclear laws, weak infrastructure and political interference may restrict it. The TOE framework therefore suggests that differences in AI-assisted supervision across Nigeria, Ghana, Kenya and South Africa arise from the interaction of technological readiness, organisational capacity and country-specific regulatory environments.

Institutional Theory extends this explanation by showing that organisational decisions respond to formal rules, professional expectations and the search for legitimacy. The theory identifies coercive, normative and mimetic pressures as important drivers of institutional behaviour. Recent research has applied these pressures to AI adoption, linking organisational decisions to regulation, professional standards and imitation of recognised adopters (Reis & Pinheiro Junior, 2025).

Coercive pressure originates from binding laws, government policies and formal accountability requirements. African financial authorities must comply with national data-protection and cybersecurity laws when processing institutional and customer information. Government digitalisation programmes and statutory responsibilities for financial stability may also compel them to develop faster and more data-driven supervision. International standards issued by the Financial Stability Board, Bank for International Settlements and other standard-setting organisations may not always have direct legal force, but they influence supervisory expectations, technical assessments and national reform priorities.

Normative pressure derives from professional values and accepted standards of sound regulatory practice. Financial regulators are expected to act independently, competently and consistently while protecting consumers and financial stability. As AI governance increasingly stresses fairness, privacy, security, explainability and accountability, these principles become professional standards against which supervisory authorities are assessed. Development institutions, regulatory networks and technical-assistance programmes also transmit norms relating to SupTech governance, model validation and responsible data use.

Mimetic pressure arises when organisations copy recognised peers during periods of uncertainty. An African regulator may introduce a supervisory sandbox, automated reporting system or AI governance policy after observing its implementation by another central bank. Such imitation can reduce experimentation costs and provide tested policy templates. It can also produce poorly suited arrangements where legal mandates, market structures, data quality and institutional resources differ. Institutional Theory thus explains why regulators may adopt visible AI initiatives to demonstrate modernity and credibility even when their technical readiness remains limited.

The theory also captures expectations from financial institutions, citizens and development organisations. Regulated firms expect clear and technologically informed oversight. Citizens expect privacy, fair treatment and protection against fraud, while governments and development partners expect effective and credible supervision. Financial authorities respond to these expectations partly to preserve institutional legitimacy. AI adoption is therefore shaped by its technical usefulness and by the regulator's need to appear capable, responsible and consistent with accepted practices.

The two theories complement each other. TOE explains whether a financial authority possesses the data, infrastructure, skills, funding and organisational readiness needed to use AI effectively. Institutional Theory explains why the authority adopts a particular approach in response to legislation, international standards, peer practices and stakeholder expectations. This combination distinguishes formal adoption from effective implementation. Institutional pressure may encourage a regulator to publish an AI strategy or launch a SupTech pilot, while weak data and limited skills prevent its meaningful use. In contrast, strong technical capability may support rapid experimentation without adequate public accountability where institutional safeguards remain weak. The combined theoretical perspective therefore supported the comparison of Nigeria, Ghana, Kenya and South Africa and informed a responsible-adoption framework that connects technical readiness, institutional capacity, regulatory legitimacy and public accountability.

2.3 Review of Empirical Studies

Recent evidence showed growing interest in AI-supported financial regulation, although direct evidence from African regulators remained limited. The World Bank (2025) examined AI adoption in emerging-market financial supervision through a survey of 27 financial authorities, including 17 African regulators, supplemented by interviews with selected authorities. The study found that financial institutions and FinTech firms were generally ahead of regulators in AI adoption. Almost half of the authorities lacked sufficient systems, resources and processes to monitor AI developments, while data privacy, cybersecurity, model transparency and dependence on third-party providers remained major concerns. Its strong African representation made it especially relevant to the present study, but the aggregation of participating authorities limited country-level comparison and did not fully disclose how supervisory approaches differed across Nigeria, Ghana, Kenya and South Africa.

Evidence on responsible governance reached similar conclusions. Bach et al. (2025) systematically reviewed 45 empirical studies involving real-world AI applications across high-risk sectors. The review identified accountability, fairness, explainability, privacy, security, AI literacy and stakeholder involvement as central responsible-AI practices. However, it found considerably more discussion of what responsible AI should involve than evidence of its practical implementation. The cross-sector coverage strengthened the governance analysis but reduced its specificity to financial supervision and African institutions. Papagiannidis et al. (2025), using a systematic review to develop a responsible-AI governance framework, similarly organised governance around structural, relational and procedural practices. Its framework clarified how responsibility could be

institutionalised, but it was not designed around the legal, infrastructural and capacity constraints of African financial regulators.

AI-assisted fraud detection represented one of the most developed application areas. Chen et al. (2025) reviewed 108 peer-reviewed studies published between 2019 and 2024 on deep-learning applications in financial fraud detection. The review found that neural-network models improved the detection of complex and evolving fraud patterns, particularly where large transaction datasets were available. It also identified persistent problems involving class imbalance, limited explainability, inconsistent datasets and weak comparability between models. The evidence supported the use of AI for suspicious-transaction analysis, but most reviewed studies measured predictive accuracy rather than supervisory usefulness, investigation outcomes or false-alert costs. Moreover, the datasets largely reflected formal digital transactions and offered little evidence on informal financial activity common in Africa.

Consumer protection and market conduct received less technical attention than fraud and prudential supervision. Fundira and Mbohwa (2025) combined a systematic review, bibliometric mapping and thematic analysis of 25 peer-reviewed studies published between 2018 and 2024. Their global review showed that AI improved service efficiency and risk management but raised concerns over discriminatory credit decisions, privacy, transparency and consumer trust. The study called for context-sensitive regulation and stronger accountability. Its limitation was the small number of reviewed publications and its emphasis on banks' use of AI rather than regulators' use of AI to detect consumer harm. Nevertheless, it established that consumer protection could not be separated from model governance and data rights.

The World Bank's 2025 evidence also showed that complaints, scams and unfair digital practices were increasingly important supervisory concerns in emerging markets. Yet authorities differed in their capacity to analyse unstructured complaint data and identify discriminatory patterns. This finding reinforced earlier work by Lopes et al. (2021), who reviewed 18 SupTech solutions used by 14 financial authorities for market-conduct supervision. The study found that automated complaint analysis, web scraping, social-media monitoring, and document analysis could improve the detection of misleading products and unfair treatment. Its case-based method offered practical examples but did not establish causal effects on consumer outcomes, and most examples came from authorities outside Africa.

Research on data privacy and cybersecurity presented a consistent caution. The Financial Stability Board (2024) conducted a global stocktake of AI adoption and related financial vulnerabilities using regulatory evidence, industry engagement, and existing research. It concluded that AI could improve operational efficiency and risk analysis but might amplify cyber risk, model risk, third-party concentration, and correlated behaviour across financial institutions. The report was broad and policy-oriented rather than an empirical country comparison, yet it showed that institution-level AI failures could acquire system-wide importance. Aldboush and Ferdous (2023), through a structured review of ethical and privacy issues in FinTech, similarly found that extensive data collection could improve personalisation and fraud controls while weakening customer trust where consent, transparency, and data-security safeguards were inadequate. Their study focused mainly on financial firms and customers, leaving the supervisory response comparatively underdeveloped.

SupTech and risk-based supervision formed another important evidence stream. Prenio (2024) analysed survey responses from 32 financial authorities and interviewed selected regulators to assess whether deployed SupTech tools had moved from experimentation into routine supervision. The study found that many tools saved time and strengthened existing supervisory processes, but only some became critical to decision-making. Management support, user involvement, data quality, appropriate skills and clear institutional strategy explained much of the difference. Because the sample was global and largely anonymised, its findings could not explain how legal mandates or capacity differences affected adoption across individual African jurisdictions. Its main relevance lay in demonstrating that acquiring or launching a tool did not amount to effective institutional adoption.

Crisanto et al. (2024) examined recent AI regulatory approaches across several jurisdictions through comparative policy and regulatory analysis. They found that authorities relied on existing rules for governance,

operational risk, data protection and model management while gradually introducing AI-specific guidance. This approach offered flexibility but created uncertainty where existing rules did not address explainability, general-purpose models and complex provider relationships. The study mainly reflected advanced and major emerging markets, limiting direct application to African regulators with different legal and institutional resources. It nevertheless provided useful dimensions for comparing whether the selected African countries applied general technology-neutral regulation, AI-specific rules or a mixture of both.

Earlier empirical evidence showed that the core implementation constraints predated generative AI. Crisanto et al. (2023) drew on responses from 55 banking, insurance and financial authorities to examine SupTech development and vendor participation. Authorities reported persistent challenges involving untimely or poor-quality data, limited staff resources and increasingly complex technology and cyber risks. Most deployed tools still concentrated on workflow automation and data architecture, while fewer used advanced AI for predictive analysis. The survey provided broad institutional evidence but did not evaluate whether the tools improved enforcement, reduced financial losses or strengthened consumer outcomes.

From a prudential perspective, Beerman et al. (2021) examined 71 SupTech tools across 20 jurisdictions. More than half analysed mainly qualitative information, confirming the value of natural-language processing for reports, supervisory assessments and other textual material. However, limited data-science skills, model calibration and data-quality problems restricted wider use. The authors also cautioned that AI outputs should support rather than replace professional judgement. Their stocktake was comprehensive for its period, but it provided little African evidence and did not address the governance questions raised by newer generative models.

Research on system-wide effects produced a less settled conclusion. Danielsson et al. (2022) used theoretical modelling and financial-market analysis to examine the connection between AI and systemic risk. They argued that AI could improve risk measurement under familiar conditions but perform poorly during rare crises for which sufficient training data did not exist. Widespread reliance on similar models could also encourage common reactions and destabilising behaviour. The study's analytical design did not test actual supervisory systems, but it challenged the assumption that greater predictive sophistication necessarily improves financial stability. This concern is directly relevant to African regulators, where incomplete historical data could make rare-event prediction even less reliable.

African evidence remained dominated by broad studies of AI development rather than financial supervision. Kondo and Diwani (2023) conducted a Scopus-based bibliometric analysis of 1,646 AI publications associated with Africa between 2013 and 2022. They found rapid publication growth but substantial geographical concentration, with South African universities leading research output. The study demonstrated expanding African technical capacity, yet its single-database design and broad disciplinary scope provided little evidence on regulatory implementation. Marak and Ayyagari (2025) later reviewed 70 Scopus-indexed studies on AI, financial inclusion and sustainable development. They concluded that AI could expand access to finance but required ethical safeguards, digital infrastructure and AI literacy. Their exclusive reliance on Scopus and focus on financial inclusion limited conclusions about supervisory institutions, although the study confirmed the importance of local infrastructure and inclusion in responsible adoption.

2.4 Synthesis of Literature and Identified Gaps

The reviewed evidence agreed that AI could improve fraud detection, regulatory-data analysis, risk prioritisation, early-warning systems and consumer monitoring. It also converged on the importance of reliable data, skilled personnel, management support, cybersecurity, explainability and clear accountability. Across the studies, effective adoption depended less on technical acquisition than on the integration of AI into supervisory routines and governance arrangements.

Disagreement remained over the extent to which AI improved supervisory effectiveness. Technical studies often reported higher predictive accuracy, while regulatory research warned that accuracy might not translate into sound decisions, fewer false alerts or better crisis management. Some studies treated automation as a

response to limited regulatory resources, whereas others showed that AI generated new demands for specialist staff, validation and monitoring. Evidence also remained mixed on whether existing technology-neutral regulation was adequate or whether AI-specific rules were required.

Fraud detection and prudential risk assessment received greater attention than market surveillance, insurance supervision, securities regulation, consumer redress and cross-border digital finance. Few studies examined contestability, the treatment of informal financial activities or responsibility for harmful supervisory decisions. African evidence was especially weak, fragmented and heavily dependent on policy reports, general AI studies and aggregated emerging-market surveys. No reviewed study offered a systematic comparison of the AI-related supervisory arrangements of Nigeria, Ghana, Kenya and South Africa across fraud detection, risk-based supervision, consumer protection, privacy, cybersecurity and institutional capacity.

These limitations established contextual, comparative, conceptual and policy gaps. Existing studies rarely combined supervisory opportunities and risks within one African analysis, and available responsible-AI frameworks generally reflected broad or advanced-market conditions. The present study addressed these gaps by systematically reviewing recent evidence, comparing four leading African financial jurisdictions and developing an African Responsible AI–Financial Supervision Framework grounded in regional legal, data, infrastructure and institutional realities.

METHODOLOGY

3.1 Research Design

The study adopted a systematic literature review integrated with comparative regulatory analysis. The systematic review provided a transparent process for locating, screening, appraising and synthesising evidence on artificial intelligence and financial-sector supervision. Comparative regulatory analysis was used to examine differences and similarities in the approaches adopted by Nigeria, Ghana, Kenya and South Africa.

The review followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses, PRISMA 2020, which supports clear reporting of the reasons for conducting a review, the search and selection procedures and the evidence finally included (Page et al., 2021). Comparative analysis was appropriate because financial regulation reflects differences in legal mandates, market structures, institutional capacity and national policy preferences (Martino, 2023). Combining the two methods allowed the study to integrate peer-reviewed research with current regulatory and institutional evidence.

3.2 Scope and Units of Analysis

The review covered publications issued between 1 January 2020 and 25 July 2026. The period captured the expansion of digital finance and SupTech following the COVID-19 pandemic, as well as recent regulatory concerns arising from generative AI.

Nigeria, Ghana, Kenya and South Africa were selected purposively. They represented major financial and FinTech markets in West, East and Southern Africa and differed in their financial-market development, legal structures, data-protection arrangements and regulatory capacities. This variation provided a suitable basis for identifying common challenges and country-specific approaches.

The units of analysis comprised peer-reviewed journal articles, credible working papers, institutional reports, legislation, policy documents and official supervisory guidance. Documents relating to banking, securities, insurance, payments, consumer protection, cybersecurity and data privacy were included where they had a clear connection to AI, SupTech or data-driven financial supervision.

3.3 Information Sources

Academic publications were retrieved from Scopus and Web of Science, while SSRN was searched for recent working papers and pre-publication research. Institutional evidence was obtained from the official repositories of the International Monetary Fund, World Bank, Bank for International Settlements and African Development Bank.

Country-level documents were retrieved from the websites of the Central Bank of Nigeria, Bank of Ghana, Central Bank of Kenya and South African Reserve Bank. Searches also covered securities regulators, deposit insurers, data-protection authorities, cybersecurity agencies and other institutions connected to financial supervision in the four countries. Reference lists of eligible documents were examined to locate additional publications not captured by the main searches.

The initial search produced 182 records from Scopus, 146 from Web of Science and 68 from SSRN. A further 87 records were obtained from the IMF, World Bank, BIS and AfDB repositories, while 64 were retrieved from African regulatory authorities. Reference searching produced another 18 records. In total, 565 records were identified.

3.4 Search Strategy

The search strategy combined three groups of terms covering AI technologies, financial supervision and geographical scope. The main Boolean search string was:

“artificial intelligence” OR “machine learning” OR “deep learning” OR “natural language processing” OR “predictive analytics” OR “SupTech” OR “supervisory technology”) AND (“financial supervision” OR “banking supervision” OR “financial regulation” OR “market conduct” OR “consumer protection” OR “fraud detection” OR “risk-based supervision”) AND (“Africa” OR “Nigeria” OR “Ghana” OR “Kenya” OR “South Africa” OR “emerging market”)

The expression was adapted to the syntax and search fields of each database. Broader searches without geographical restrictions captured international evidence on responsible AI, model risk, explainability and SupTech. Country-specific searches combined each country’s name with terms such as “central bank,” “AI strategy,” “regulatory sandbox,” “data protection,” “cybersecurity,” “consumer protection” and “digital finance regulation.”

For institutional websites with limited search functions, shorter combinations were used. The database, search date, exact search expression, filters and number of records returned were entered into a search register to support transparency and reproducibility.

3.5 Eligibility Criteria

Publications were included where they:

- were published between January 2020 and July 2026;
- addressed AI, machine learning, SupTech or data-driven financial supervision;
- examined fraud detection, risk-based supervision, regulatory reporting, consumer protection, data privacy, cybersecurity, market surveillance or institutional capacity;
- were peer-reviewed articles, credible working papers, institutional reports, legislation or official guidance;
- were available in English; and
- contained adequate methodological, analytical or policy information.

The review excluded purely technical studies that evaluated algorithms without discussing regulatory or supervisory relevance. General business-automation studies, commercial publications, unsupported opinion pieces, duplicate records and documents outside the review period were also excluded. Where a working paper and a later journal article reported substantially the same evidence, the journal version was retained unless the earlier paper contained unique regulatory information.

3.6 PRISMA Selection Procedure

The selection process followed identification, deduplication, screening, eligibility assessment, and final inclusion. Of the 565 records identified, 141 duplicates were removed, leaving 424 records for title and abstract screening. A total of 296 records were excluded at this stage because they lacked financial-supervision relevance, addressed general AI applications, or failed to meet the period and publication criteria.

The full texts of 128 documents were sought. Eight could not be retrieved, leaving 120 documents for eligibility and quality assessment. Sixty-four full texts were excluded: 22 lacked direct supervisory relevance, 15 had a purely technical focus, 12 contained insufficient methodological or policy evidence, nine duplicated evidence available elsewhere and six fell outside the review period. The final analytical sample therefore comprised 56 publications.

The sample contained 28 peer-reviewed journal articles, seven working papers, 14 institutional or policy reports and seven laws, regulations or official supervisory documents. The inclusion of both academic and institutional sources was necessary because regulatory developments often appeared in official reports before they received extensive peer-reviewed examination.

3.7 Quality Appraisal

Each eligible publication was assessed against six criteria: clarity of purpose, appropriateness of method, credibility of evidence, transparency of data sources, relevance to financial supervision and strength of conclusions. Each criterion received 0 where it was absent, 1 where it was partly satisfied and 2 where it was fully satisfied. Total scores ranged from 0 to 12.

Publications scoring 9–12 were classified as high quality, those scoring 6–8 as moderate quality and those below 6 as low quality. Thirty-five publications, representing 62.5% of the sample, were rated as high quality, while 21 publications, representing 37.5%, were rated as moderate quality. Low-quality publications were excluded from the main synthesis.

Laws and official guidance were assessed by legal status, institutional authority, currency, clarity and relevance rather than solely by conventional empirical research criteria. This distinction prevented regulatory documents from being assessed using standards designed exclusively for journal articles.

3.8 Data Extraction

A standard extraction matrix was used to record the author and year, publication type, country or region, research method, supervisory function and AI application examined. The matrix also captured identified opportunities, risks, regulatory safeguards, institutional constraints, principal findings and quality ratings.

The use of a uniform extraction form reduced inconsistency and allowed evidence from different publication types to be compared. Where a document addressed several themes, one principal theme was assigned for descriptive classification, while secondary themes were retained during qualitative coding.

3.9 Comparative Regulatory Analysis

The country analysis applied a common matrix covering eight dimensions:

1. AI or digital-finance policy;

2. SupTech initiatives;
3. data-protection framework;
4. cybersecurity requirements;
5. consumer-protection arrangements;
6. regulatory sandboxes;
7. institutional capacity; and
8. accountability and oversight.

Each dimension was coded on a three-point scale. A score of 2 indicated an established law, policy, operational programme or responsible institution. A score of 1 indicated a developing arrangement, such as a draft policy, pilot programme or partial institutional framework. A score of 0 indicated that no relevant arrangement was identified in the reviewed public documents.

The analysis considered the legal status, scope, implementation evidence and institutional responsibility attached to each arrangement. The existence of a general AI strategy or sandbox was not treated automatically as evidence that AI had become embedded in financial supervision. This prevented formal policy adoption from being confused with operational supervisory capacity.

3.10 Method of Analysis

The selected documents were analysed using thematic content analysis. The process involved familiarisation with the publications, initial coding, consolidation of related codes into themes, comparison across countries and identification of policy gaps. The six primary themes were fraud detection, risk-based supervision, consumer protection, data privacy, cybersecurity and institutional capacity.

Deductive and inductive coding were combined. The deductive categories reflected the study's research questions, while additional patterns arising from the documents were retained where they improved the analysis. Contradictory findings, implementation failures and missing evidence were examined alongside favourable accounts of AI adoption.

Consistent with Braun and Clarke (2022), themes were developed as meaningful analytical patterns rather than simple counts of repeated words. The thematic findings and regulatory comparison were subsequently integrated into the African Responsible AI–Financial Supervision Framework. Framework construction followed a theory-to-evidence mapping procedure. Technology–Organisation–Environment (TOE) categories were matched to the technological, organisational and environmental constraints identified in the review, while Institutional Theory was used to interpret legal obligations, professional standards, peer diffusion and legitimacy pressures. Each proposed pillar was retained only where it responded to a recurrent regulatory challenge and could be linked to a defined governance practice and supervisory outcome.

3.11 Trustworthiness and Reliability

Several procedures strengthened the reliability of the review. These included a documented search protocol, consistent eligibility criteria, a standard extraction matrix and an audit trail of screening and exclusion decisions. Publication classifications and quality scores were cross-checked to identify coding inconsistencies.

The study combined peer-reviewed and institutional evidence to balance methodological quality with current regulatory information. It also retained negative and conflicting evidence to reduce selective reporting. The search strings, excluded full texts, quality scores and country-coding decisions were documented to permit verification.

3.12 Ethical Considerations

The study relied exclusively on publicly available documents and did not involve human participants or collect personal data. Formal participant consent was therefore not required. Ethical responsibility remained necessary in the accurate representation of sources, disclosure of screening decisions and avoidance of selective reporting.

The absence of publicly available evidence was not interpreted automatically as proof that a regulator had undertaken no relevant activity. Regulatory documents were also interpreted according to their dates, legal status and institutional context. This approach reduced the risk of overstating policy implementation or misrepresenting evolving national arrangements.

FINDINGS AND DISCUSSION

4.1 PRISMA Search and Selection Results

The systematic search identified 565 records from academic databases, institutional repositories, African regulatory authorities and reference searches. As presented in Table 1, Scopus produced the highest number of records, followed by Web of Science and institutional repositories. After removing 141 duplicates, 424 records remained for title and abstract screening.

Table 1: PRISMA Search and Selection Results

PRISMA stage	Activity	Number
Identification	Records retrieved from Scopus	182
	Records retrieved from Web of Science	146
	Records retrieved from SSRN	68
	Records retrieved from IMF, World Bank, BIS and AfDB	87
	Records retrieved from African regulatory authorities	64
	Records identified through reference searches	18
	Total records identified	565
Deduplication	Duplicate records removed	141
Screening	Records screened by title and abstract	424
	Records excluded during initial screening	296
Retrieval	Full-text documents sought	128
	Full-text documents not retrieved	8
Eligibility	Full-text documents assessed	120
	Full-text documents excluded	64

Inclusion	Publications included in the final review	56
-----------	---	----

Of the 424 screened records, 296 were removed because they lacked financial-supervision relevance, addressed general AI applications or failed to meet the eligibility criteria. The full texts of 128 documents were sought, of which eight could not be retrieved. The remaining 120 documents underwent full-text and quality assessment. Sixty-four were excluded, leaving 56 publications in the final synthesis.




Figure 2. PRISMA flow diagram for the identification and selection of publications.

Source: Authors' construction based on Page et al. (2021).

Consistency check

- Records identified: $396 + 169 = 565$ $396 + 169 = 565$ $396 + 169 = 565$

- Records screened: $565 - 141 = 424$ $565 - 141 = 424$
- Full texts sought: $424 - 296 = 128$ $424 - 296 = 128$
- Full texts assessed: $128 - 8 = 120$ $128 - 8 = 120$
- Final publications: $120 - 64 = 56$ $120 - 64 = 56$
- Publication types: $28 + 7 + 14 + 7 = 56$ $28 + 7 + 14 + 7 = 56$

Table 2: Reasons for Full-Text Exclusion

Reason for exclusion	Number	Percentage
No direct financial-supervision relevance	22	34.4
Purely technical AI study	15	23.4
Insufficient methodological or policy evidence	12	18.8
Duplicate or substantially overlapping evidence	9	14.1
Outside the review period	6	9.4
Total	64	100

Lack of direct supervisory relevance was the leading reason for exclusion. This result reflected a broader weakness in the literature: many studies evaluated AI models for credit, fraud or forecasting without examining how regulators could validate, supervise or apply those models. The selection process therefore retained publications with clear regulatory, governance or supervisory relevance.

4.2 Profile of Included Publications

Publication activity increased during the review period, particularly from 2023 onwards. As shown in Table 3, 32 publications, representing 57.1% of the final sample, were published between 2024 and 2026. The concentration in recent years reflected increased regulatory attention to generative AI, model risk, cybersecurity and responsible adoption.

Table 3: Distribution of Included Publications by Year

Publication year	Frequency	Percentage
2020	3	5.4
2021	5	8.9
2022	7	12.5
2023	9	16.1
2024	12	21.4
2025	15	26.8

2026	5	8.9
Total	56	100

Peer-reviewed journal articles constituted 50.0% of the sample. Institutional and policy reports accounted for another 25.0%, confirming the prominent role of the World Bank, BIS, IMF, AfDB and financial authorities in documenting SupTech developments.

Table 4: Distribution by Publication Type

Publication type	Frequency	Percentage
Peer-reviewed journal articles	28	50
Working papers	7	12.5
Institutional and policy reports	14	25
Laws, regulations and official supervisory guidance	7	12.5
Total	56	100

The geographical distribution remained uneven. Global and multi-country studies represented 33.9% of the sample, while only 13 publications focused directly on Nigeria, Ghana, Kenya or South Africa.

Table 5: Geographical Coverage of Included Publications

Geographical coverage	Frequency	Percentage
Global or multi-country studies	19	33.9
Advanced economies	8	14.3
Emerging markets and developing economies	9	16.1
Africa generally	7	12.5
Nigeria	3	5.4
Ghana	2	3.6
Kenya	3	5.4
South Africa	5	8.9
Total	56	100

Note: Percentages may differ slightly from 100% because of rounding.

South Africa attracted the greatest country-specific coverage, while Ghana received the least. This finding was consistent with Kondo and Diwani's (2023) observation that African AI research was expanding but remained geographically concentrated around South African institutions.

Comparative regulatory analysis and systematic reviews were the dominant methods, jointly accounting for 50.0% of the sample. Only eight studies used surveys or interviews, while three adopted mixed methods.

Table 6: Distribution by Research Method

Research method	Frequency	Percentage
Comparative regulatory or document analysis	15	26.8
Systematic or structured literature review	13	23.2
Survey or interview study	8	14.3
Technical model evaluation	7	12.5
Case study	6	10.7
Conceptual or analytical modelling	4	7.1
Mixed-method study	3	5.4
Total	56	100

The dominance of documentary methods indicated that the field remained at an early stage. Few studies measured whether AI-assisted supervision reduced financial losses, accelerated intervention or improved consumer outcomes.

Fraud detection and prudential supervision received the greatest attention. Securities-market surveillance and institutional capacity were comparatively under-researched.

Table 7: Distribution by Principal Supervisory Theme

Principal supervisory theme	Frequency	Percentage
Fraud and financial-crime detection	10	17.9
Prudential and risk-based supervision	9	16.1
Data privacy and cybersecurity	7	12.5
Responsible-AI governance	7	12.5
Regulatory reporting and data analysis	6	10.7
Consumer protection and market conduct	6	10.7
Early-warning and financial-stability monitoring	5	8.9
Institutional capacity	4	7.1
Securities-market surveillance	2	3.6
Total	56	100

Thirty-five publications were rated as high quality, while 21 satisfied the moderate-quality threshold. Publications classified as low quality were excluded from the main synthesis.

Table 8: Quality Classification of Included Publications

Quality classification	Score range	Frequency	Percentage
High quality	9–12	35	62.5
Moderate quality	6–8	21	37.5
Low quality	Below 6	0	0
Total		56	100

Country-level evidence depended heavily on official documents and institutional reports. Only four peer-reviewed studies focused directly on the selected countries.

Table 9: Country-Specific Publications by Type

Country	Peer-reviewed studies	Working papers	Institutional reports	Official documents	Total
Nigeria	1	0	1	1	3
Ghana	0	0	1	1	2
Kenya	1	0	1	1	3
South Africa	2	0	1	2	5
Total	4	0	4	5	13

The profile confirmed a central research gap. Global discussions of financial AI had expanded faster than country-specific evidence on how African authorities adopted, governed and evaluated AI-assisted supervision.

4.3 AI Applications in Financial-Sector Supervision

The reviewed publications identified actual, emerging and proposed AI applications. Actual uses were concentrated in regulatory-return validation, document classification, complaint analysis, risk dashboards and transaction monitoring. These applications largely automated existing supervisory processes rather than replacing established supervisory models.

Natural-language processing was used to examine regulatory returns, inspection reports, audit findings, complaints and other unstructured information. Machine-learning systems compared institutions, identified unusual movements and supported supervisory risk classification. Beerman et al. (2021), in their review of 71 SupTech tools across 20 jurisdictions, found that more than half primarily analysed qualitative information. This finding demonstrated the value of text-based AI where financial supervision depended heavily on reports, correspondence and supervisory assessments.

Emerging applications included predictive risk scoring, transaction-network analysis, automated anomaly detection and early-warning models. These systems combined capital, liquidity, profitability, governance and market indicators to identify institutions requiring closer attention. AI also assisted fraud detection by identifying unusual payment patterns and relationships that rule-based systems might overlook. Chen et al. (2025) found that deep-learning methods improved the detection of complex financial fraud, although differences in datasets and evaluation procedures limited direct comparison between models.

Proposed uses included generative-AI assistants for regulatory research, summarisation of inspection documents and conversational access to supervisory databases. These applications could reduce the time spent reviewing large volumes of material. However, the evidence did not support their autonomous use for licensing, enforcement or high-impact risk classification. Prenio (2024) found that deployed SupTech systems added the most value when they strengthened existing supervisory work and remained subject to professional judgement. BIS evidence on SupTech deployment

4.4 Opportunities for African Financial Authorities

Fraud and financial-crime detection represented the most immediate opportunity. African financial systems have experienced rapid growth in electronic payments, mobile money and digital credit, increasing both transaction volumes and exposure to technology-enabled fraud. AI can identify abnormal behaviour, trace linked transactions and prioritise alerts for investigation. Its practical benefit, however, depends on data quality and the capacity of enforcement teams to investigate flagged cases.

AI also strengthens risk-based supervision. By combining regulatory returns, financial ratios, governance information and market indicators, supervisory authorities can direct scarce resources towards institutions presenting the greatest potential risk. Dohotaru et al. (2025) found that AI could support risk-based supervision in credit risk, stress testing, consumer protection and anti-money-laundering activities. This capability is relevant where financial-market expansion has outpaced staffing and supervisory budgets.

Early-warning systems can flag deteriorating asset quality, unusual liquidity movements, rising complaints or abnormal payment behaviour before conventional reporting cycles reveal them. Nevertheless, such alerts remain probabilistic. Danielsson et al. (2022) cautioned that AI models trained on normal market conditions might perform poorly during rare crises and could increase systemic risk where institutions relied on similar models.

Consumer protection offered another important application. Natural-language processing can organise complaints, analyse product conditions and identify recurring allegations of hidden charges, discriminatory decisions or aggressive recovery practices. Lopes et al. (2021) documented the use of complaint analytics, web scraping and document analysis for market-conduct supervision. Yet only six publications in the present review concentrated principally on consumer protection, indicating that this supervisory function remained less developed than fraud and prudential analysis.

AI also improved regulatory efficiency by automating repetitive validation, classification and reporting tasks. This could allow professional staff to concentrate on investigation, judgement and enforcement. The efficiency benefit should not be interpreted as reducing the need for supervisors. AI created new requirements for data management, model validation, cybersecurity and accountability.

4.5 Risks and Implementation Constraints

Poor data quality was the most consistent constraint. Fragmented databases, inconsistent reporting formats, incomplete histories and limited coverage of informal activities could produce misleading supervisory classifications. An apparently accurate system might still exclude rural institutions, informal providers or consumers with limited digital records.

Model risk arose through faulty assumptions, overfitting, inappropriate application and performance deterioration. Historical data could reproduce existing inequalities, creating biased results across gender, location, income or business type. Fundira and Mbohwa (2025) identified bias, privacy, transparency and accountability as persistent concerns in AI-enabled banking. Limited explainability further weakened the ability of supervisors and regulated entities to understand or challenge consequential outputs.

Privacy risks increased as regulators acquired more granular customer and transaction data. Centralised datasets also created valuable targets for cyberattack. AI systems could be manipulated through altered inputs,

compromised training data or attacks designed to evade fraud controls. The Financial Stability Board (2024) warned that AI could intensify cyber risk, model risk and dependence on a small number of cloud, hardware and model providers.

Accountability remained difficult where AI influenced inspection, licensing or enforcement decisions. A named official or regulatory body must retain responsibility for final decisions, while affected parties need a process for correction and review. Bach et al. (2025) found broad support for accountability, fairness, explainability and stakeholder participation but limited evidence that these principles were consistently implemented.

African regulators also faced shortages of data scientists, cybersecurity specialists and model validators. Funding limitations, weak interoperability and reliance on external vendors increased these constraints. Consequently, poorly planned adoption could automate existing weaknesses rather than improve supervision.

4.6 Comparative Regulatory Findings

The four countries had made progress in data protection, cybersecurity, consumer protection and controlled financial innovation. However, their readiness for AI-enabled supervision differed.

Table 10: Comparative Regulatory Position of the Selected Countries

Dimension	Nigeria	Ghana	Kenya	South Africa
AI or digital-finance policy	Developing	Developing	Established national AI strategy	Established sectoral assessment
Disclosed SupTech initiatives	Developing	Developing	Developing	Relatively established
Data-protection framework	Established	Established	Established	Established
Cybersecurity requirements	Established	Established	Established	Established
Consumer protection	Established	Established	Established	Established
Regulatory sandbox	Operational	Operational	Operational	Operational and inter-regulatory
Institutional capacity	Developing	Developing	Developing	Relatively established
AI accountability structure	Fragmented	Fragmented	Developing	More clearly coordinated

Nigeria had a National AI Strategy, data-protection legislation, CBN cybersecurity supervision and operational innovation arrangements. However, there was limited public evidence that advanced AI systems had become part of routine financial supervision. AI-related responsibilities were also distributed across several agencies.

Ghana had established data-protection and cybersecurity requirements and operated a regulatory sandbox. Its approach supported controlled financial innovation, but public evidence of AI-specific SupTech applications and a unified accountability structure remained limited.

Kenya had the clearest national AI strategy, supported by considerable experience in mobile finance, digital-credit regulation and regulatory sandboxes. Its main challenge was converting broad national AI principles into specific and enforceable standards for supervised financial institutions.

South Africa demonstrated the strongest disclosed readiness. The Prudential Authority and Financial Sector Conduct Authority had examined AI applications and risks in the financial sector, while the Twin Peaks framework clarified prudential and conduct responsibilities. Its remaining challenges included explainability, cyber risk, third-party dependence and specialist capacity.

Transferable practices included Nigeria's cybersecurity assessment, Ghana's innovation sandbox, Kenya's national AI strategy and South Africa's joint prudential-conduct assessment. Regional coordination remained necessary for cross-border data sharing, fraud intelligence, cloud-provider risk, model validation and oversight of digital platforms operating across several countries.

4.7 Discussion in Relation to the Theories

The findings supported the Technology–Organisation–Environment framework. Technological conditions, particularly data quality, interoperability, cybersecurity and infrastructure, determined whether AI applications could produce reliable supervisory information. Organisational factors such as leadership support, funding, specialist expertise and user involvement affected whether pilot systems became embedded in routine supervision. Environmental conditions, including legislation, market complexity, political priorities and international standards, influenced both the pace and form of adoption.

The country comparison also supported Institutional Theory. Data-protection laws, cybersecurity standards and government digitalisation policies created coercive pressure. Guidance from the BIS, World Bank, FSB and other international bodies generated normative pressure by defining accepted approaches to AI governance. The spread of regulatory sandboxes and risk-based AI principles reflected mimetic pressure, as authorities adopted practices observed in peer jurisdictions.

Institutional similarity did not imply equal implementation. An authority could launch a sandbox or publish an AI strategy to signal regulatory modernity without possessing the data, skills or infrastructure required for effective supervision. TOE therefore explained adoption capacity, while Institutional Theory explained the external pressures and legitimacy concerns shaping regulatory choices. Their combined use distinguished formal policy adoption from operational readiness.

4.8 African Responsible AI–Financial Supervision Framework

The evidence supported a framework comprising six connected pillars. The framework is theoretically grounded in TOE and Institutional Theory rather than presented as a stand-alone normative model. TOE provides the readiness logic: its technological context informs data governance and secure, explainable systems; its organisational context informs institutional capacity, accountability and monitoring; and its environmental context informs legal foundations, proportionality and inter-agency cooperation. Institutional Theory provides the governance and legitimacy logic. Coercive pressures explain the need for lawful mandates, privacy duties and enforceable accountability; normative pressures support professional standards for fairness, validation and auditability; and mimetic pressures justify controlled experimentation and regional learning while guarding against copying models that do not fit local conditions.

The pillars also answer the specific weaknesses identified in the review. Legal and ethical foundations address unclear mandates, fragmented rules and risks to fairness. Data governance and privacy respond to incomplete, poor-quality and unlawfully shared data. Secure and explainable systems address cyber threats, opaque models, unreliable outputs and vendor dependence. Proportionate supervisory use limits harm where AI affects licensing, risk ratings, sanctions or consumer outcomes. Institutional capacity and cooperation respond to skills shortages, weak infrastructure and divided regulatory mandates. Accountability, monitoring and review

address responsibility gaps, model deterioration, limited contestability and the danger that automated outputs displace documented supervisory judgement.




Figure 3: Institutional Framework and Phased Adoption Cycle for Responsible Artificial Intelligence in African Financial Institutions

Source: Author's conceptualisation based on the reviewed literature and regulatory frameworks (2026).

Table 11: African Responsible AI–Financial Supervision Framework

Pillar	Regulatory challenge addressed	Governance response and expected result
Legal and ethical foundations	Unclear statutory authority; fragmented AI rules; risks of unfair or disproportionate action	Define lawful purpose, approved and prohibited uses, fairness duties and proportionality tests, producing legitimate use consistent with supervisory mandates.

Data governance and privacy	Incomplete and inconsistent data; weak provenance; unlawful access or sharing; privacy breaches	Apply quality standards, provenance records, minimisation, access controls and secure sharing, producing more reliable analysis and stronger protection of sensitive information.
Secure and explainable systems	Cyberattack, model manipulation, opaque outputs, weak validation and dependence on external vendors	Require security testing, model documentation, validation, interpretable reasons, audit rights and continuity plans, producing defensible outputs and lower technical vulnerability.
Proportionate supervisory use	Uniform controls that ignore differences between low-risk analytics and consequential licensing, sanctions or risk ratings	Match approval, review and intervention controls to the consequence of each use, with stronger safeguards and documented judgement for high-impact decisions.
Institutional capacity and cooperation	Shortages of specialist skills and funding; weak infrastructure; fragmented mandates and cross-border risks	Build multidisciplinary teams, shared infrastructure, training and inter-agency or regional arrangements, supporting sustainable implementation and coordinated oversight.
Accountability, monitoring and review	Unclear responsibility, model drift, weak contestability, limited incident learning and overreliance on automated outputs	Assign named responsibility, maintain audit trails, provide challenge and remedy procedures, monitor performance and require independent assessment, preserving continuing control over AI-supported decisions.

The framework operates through the following cycle:

Readiness assessment → controlled experimentation → validation → phased deployment → continuous monitoring → independent review

Readiness assessment determines whether the authority has a lawful mandate, suitable data, adequate infrastructure and necessary skills. Controlled experimentation permits testing in a restricted environment. Validation examines accuracy, bias, security, explainability and operational relevance. Phased deployment limits initial use to lower-risk functions before expansion. Continuous monitoring detects model deterioration and emerging threats, while independent review confirms continued legality, proportionality and effectiveness.

The cycle is not strictly linear. Failure at any stage requires the application to return to an earlier stage for correction. Responsible adoption is therefore treated as an ongoing supervisory process rather than a one-time technology acquisition.

4.9 Broader Policy and Regulatory Implications

African central banks should establish inventories of AI systems used internally and by regulated financial institutions. Minimum standards should cover model documentation, data provenance, validation, incident reporting, cybersecurity and third-party providers. AI-supported alerts should remain subject to documented supervisory judgement, particularly where they influence licensing, sanctions or institutional risk ratings.

Securities regulators require greater capacity for automated market surveillance and oversight of algorithmic trading. Insurance regulators need sector-specific tools for pricing, underwriting, claims and consumer-conduct

risks. Deposit insurers can use AI-supported monitoring to strengthen early identification of institutional distress, provided model outputs are independently validated.

Data-protection authorities should work directly with financial regulators on lawful data sharing, automated decisions and consumer remedies. Cybersecurity agencies should support common incident classifications and financial-sector threat intelligence. These arrangements are necessary because AI-related risks rarely fall within one institution's mandate.

The African Union, AfDB and regional economic communities can support shared training, model-validation facilities and minimum responsible-AI standards. Regional arrangements are especially important for digital lenders, payment providers, fraud networks and technology vendors operating across national borders.

Overall, the findings supported gradual and risk-based AI adoption. AI offered clear opportunities for fraud detection, risk prioritisation, early-warning analysis and regulatory efficiency, but its benefits depended on reliable data, institutional capacity and enforceable safeguards. The evidence supported AI as a tool for strengthening professional supervision, not as a substitute for accountable regulatory judgement.

SUMMARY, RECOMMENDATIONS AND CONCLUSION

Summary of Major Findings

This study systematically reviewed 56 academic and regulatory publications and compared AI-related financial-supervision arrangements in Nigeria, Ghana, Kenya and South Africa. The analysis addressed five research objectives.

First, the findings showed that AI applications in financial supervision were expanding, although most African authorities remained at an early stage of operational adoption. Current applications centred on regulatory-return validation, fraud monitoring, complaint classification, document analysis and supervisory risk dashboards. Emerging uses included predictive risk scoring, natural-language processing, transaction-network analysis and early-warning systems. Generative-AI applications for regulatory research and document summarisation remained largely experimental.

Second, AI offered substantial opportunities for supervisory effectiveness. Fraud detection received the greatest research attention, accounting for 17.9% of the reviewed publications, followed by prudential and risk-based supervision at 16.1%. AI could process large transaction volumes, identify unusual behaviour, improve regulatory reporting and direct limited supervisory resources towards higher-risk institutions. It could also support consumer protection by analysing complaints, advertisements and digital-credit practices. However, evidence of measurable improvements in regulatory intervention, consumer redress and financial stability remained limited.

Third, the review identified serious risks and institutional constraints. Poor data quality, fragmented databases and the limited coverage of informal financial activities could weaken model reliability. Algorithmic bias, inadequate explainability, privacy breaches, cyberattacks and dependence on external technology providers also created regulatory concerns. These risks were compounded by shortages of data scientists, cybersecurity specialists and model validators. Accountability remained especially important where AI outputs influenced inspection, licensing or enforcement.

Fourth, the country comparison showed that the four jurisdictions had established important foundations through data-protection laws, cybersecurity requirements, consumer-protection arrangements and regulatory sandboxes. Kenya had the clearest national AI strategy, while South Africa demonstrated the most developed financial-sector AI assessment and institutional readiness. Nigeria had strengthened cybersecurity assessment and digital-finance regulation, whereas Ghana had developed an operational innovation sandbox and formal cybersecurity requirements. However, all four countries had limited publicly available evidence showing that advanced AI had become embedded in routine supervision.

Fifth, the findings supported an African Responsible AI–Financial Supervision Framework built around six pillars: legal and ethical foundations, data governance and privacy, secure and explainable systems, proportionate supervisory use, institutional capacity and cooperation, and accountability, monitoring and review. These pillars operated through a continuous adoption cycle comprising readiness assessment, controlled experimentation, validation, phased deployment, continuous monitoring and independent review.

Recommendations

National financial regulators should adopt formal AI-governance policies covering both their internal systems and AI applications used by regulated institutions. These policies should define approved uses, prohibited practices, responsibility for decisions, documentation requirements and procedures for challenging consequential outputs. Regulators should maintain registers of high-risk AI systems and require periodic reporting on model performance, bias, cyber incidents and third-party dependence.

African central banks should expand supervisory sandboxes beyond product innovation to include controlled testing of SupTech tools. Pilot applications should initially focus on lower-risk activities such as return validation, complaint classification and document analysis. Systems used for institutional risk ratings, licensing or enforcement should face stronger validation, approval and independent review. Central banks should also establish multidisciplinary teams combining supervisory, legal, data-science, cybersecurity and consumer-protection expertise.

Data-protection and cybersecurity authorities should work with financial regulators to develop common standards for lawful data access, automated decisions, cross-border data transfers and incident reporting. Material AI-related failures, including discriminatory outcomes, data breaches, manipulated models and extended system outages, should be reported within clearly defined periods. Joint inspection and information-sharing procedures would reduce fragmented oversight.

Financial institutions and technology providers should maintain complete records of training data, model assumptions, performance tests, changes and responsible personnel. Vendors should provide regulators with sufficient access to assess proprietary systems. Contracts should address audit rights, data location, cybersecurity responsibilities, service continuity and exit arrangements. Institutions should not rely solely on vendor assurances when the system influences credit, customer treatment, fraud controls or regulatory reporting.

The AfDB, African Union and regional economic communities should support shared regulatory infrastructure for countries unable to develop advanced SupTech independently. Regional initiatives could include secure fraud-intelligence platforms, common data dictionaries, model-validation laboratories and specialist training programmes. Minimum regional principles for fairness, explainability, privacy, cybersecurity and accountability would reduce regulatory gaps while allowing national adaptation.

Research institutions, universities and professional bodies should develop practical training in supervisory analytics, model-risk management, AI auditing and financial-sector cybersecurity. Partnerships with regulators should produce African datasets, case studies and validation methods. Independent researchers should also participate in periodic assessment of high-impact regulatory systems to strengthen transparency and public confidence.

Conclusion

AI can strengthen financial-sector supervision in Africa by improving fraud detection, risk prioritisation, regulatory reporting, consumer monitoring and early identification of financial vulnerabilities. These benefits are particularly relevant where supervisory responsibilities are increasing faster than available human and financial resources.

However, the effectiveness of AI depends on the environment in which it operates. Models trained on incomplete or unrepresentative data may automate existing weaknesses and exclusion. Limited explainability can weaken accountability, while inadequate cybersecurity and excessive vendor dependence can create new sources of operational and systemic risk. Responsible adoption must therefore reflect local data conditions, legal protections, institutional capacity and financial-market structures.

The central issue is not whether African regulators should adopt AI, but how they can do so without weakening fairness, public accountability and supervisory judgement. Gradual deployment, clear governance, reliable data, skilled personnel and independent review provide the most credible path. AI should support informed regulatory decisions rather than replace the public officials responsible for them.

Limitations and Future Research

The study depended on published and publicly available evidence and did not collect primary data from regulators, financial institutions or FinTech firms. It therefore assessed disclosed regulatory readiness rather than the full extent of operational readiness or day-to-day governance practice. Undisclosed internal projects, unsuccessful pilots, confidential supervisory systems and informal implementation barriers may have been omitted. Uneven disclosure across the four countries also made it difficult to distinguish the absence of an initiative from the absence of public information. Rapid changes in AI technology and regulation may reduce the currency of some findings, while limited country-specific peer-reviewed evidence restricted assessment of actual supervisory outcomes. The framework should therefore be treated as an analytically derived and context-sensitive proposition that requires empirical validation, rather than as a tested causal model.

Future research should test the framework with primary data from central banks and other financial regulators, deposit money banks, insurers, securities firms, payment-service providers, FinTech companies, data-protection authorities and cybersecurity agencies. Semi-structured interviews could examine implementation barriers, allocation of responsibility, model-validation practice, vendor oversight, staff capability and inter-agency coordination. A complementary survey could measure institutional readiness across the six pillars using indicators such as legal mandate clarity, data quality, secure infrastructure, specialist staffing, governance maturity and independent-review capacity. Comparative case studies or pilot implementations in selected African financial systems should then trace how the pillars operate in practice, including their data requirements, accuracy, false-alert rates, bias, explainability, incident response and consumer-remedy arrangements. Longitudinal and mixed-method studies could determine whether framework adoption improves supervisory intervention speed, reduces fraud losses, strengthens consumer protection and produces sustained institutional capability. Such testing would permit refinement of the pillar measures, assessment of relationships proposed by TOE and Institutional Theory, and evaluation of the framework's transferability beyond Nigeria, Ghana, Kenya and South Africa.

REFERENCES

1. Aldboush, H. H. H., & Ferdous, M. (2023). Building trust in fintech: An analysis of ethical and privacy considerations in the intersection of big data, AI, and customer trust. *International Journal of Financial Studies*, 11(3), Article 90. <https://doi.org/10.3390/ijfs11030090>
2. Alshahrani, A., Dennehy, D., & Mäntymäki, M. (2022). An attention-based view of AI assimilation in public sector organizations: The case of Saudi Arabia. *Government Information Quarterly*, 39(4), Article 101617. <https://doi.org/10.1016/j.giq.2021.101617>
3. Bach, T. A., Kaarstad, M., Solberg, E., & Babic, A. (2025). Insights into suggested responsible AI practices in real-world settings: A systematic literature review. *AI and Ethics*, 5, 3185–3232. <https://doi.org/10.1007/s43681-024-00648-7>
4. Bahoo, S., Cucculelli, M., Goga, X., & Mondolo, J. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. *SN Business & Economics*, 4, Article 23. <https://doi.org/10.1007/s43546-023-00618-x>

5. Beerman, K., Prenio, J., & Zamil, R. (2021). Suptech tools for prudential supervision and their use during the pandemic (FSI Insights No. 37). Bank for International Settlements. <https://www.bis.org/fsi/publ/insights37.htm>
6. Braun, V., & Clarke, V. (2023). Toward good practice in thematic analysis: Avoiding common problems and be(com)ing a knowing researcher. *International Journal of Transgender Health*, 24(1), 1–6. <https://doi.org/10.1080/26895269.2022.2129597>
7. Chen, Y., Zhao, C., Xu, Y., Nie, C., & Zhang, Y. (2025). Deep learning in financial fraud detection: Innovations, challenges, and applications. *Data Science and Management*. Advance online publication. <https://doi.org/10.1016/j.dsm.2025.08.002>
8. Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024). Regulating AI in the financial sector: Recent developments and main challenges (FSI Insights No. 63). Bank for International Settlements. <https://www.bis.org/fsi/publ/insights63.htm>
9. Danielsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140, Article 106290. <https://doi.org/10.1016/j.jbankfin.2021.106290>
10. Dohotaru, M., Prisacaru, M., Shin, J. H., & Palta, Y. (2025). AI for risk-based supervision: Another “nice to have” tool or a game-changer? World Bank. <https://openknowledge.worldbank.org/entities/publication/4f8b828d-3eb0-4f16-a816-b5143c414e70>
11. Financial Stability Board. (2024). The financial stability implications of artificial intelligence. <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>
12. Fundira, M., & Mbohwa, C. (2025). AI ethics in banking services: A systematic and bibliometric review of regulatory and consumer perspectives. *Discover Artificial Intelligence*, 5, Article 319. <https://doi.org/10.1007/s44163-025-00432-4>
13. Khanfar, A. A., Kiani Mavi, R., Iranmanesh, M., & Gengatharen, D. (2026). Determinants of artificial intelligence adoption: Research themes and future directions. *Information Technology and Management*, 27, 31–51. <https://doi.org/10.1007/s10799-024-00435-0>
14. Kondo, T. S., & Diwani, S. A. (2023). Artificial intelligence in Africa: A bibliometric analysis from 2013 to 2022. *Discover Artificial Intelligence*, 3, Article 34. <https://doi.org/10.1007/s44163-023-00084-2>
15. Lopes, L., Chien, J., Wallace, M., & Totolo, E. (2021). The next wave of suptech innovation: Suptech solutions for market conduct supervision. World Bank. <https://openknowledge.worldbank.org/entities/publication/79d2a5dc-c88b-58ee-834e-a175f7b165b6>
16. Marak, N. R., & Ayyagari, L. R. (2025). Artificial intelligence for financial inclusion and sustainable development: A systematic literature review. *Discover Artificial Intelligence*, 5, Article 390. <https://doi.org/10.1007/s44163-025-00668-0>
17. Martino, E. D., Nabilou, H., & Paccès, A. M. (2023). Comparative financial regulation: The analytical framework (ECGI Law Working Paper No. 726/2023). European Corporate Governance Institute. <https://doi.org/10.2139/ssrn.4522348>
18. Organisation for Economic Co-operation and Development. (2024a). Explanatory memorandum on the updated OECD definition of an AI system (OECD Artificial Intelligence Papers No. 8). OECD Publishing. <https://doi.org/10.1787/623da898-en>
19. Organisation for Economic Co-operation and Development. (2024b). Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449). <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>
20. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, Article n71. <https://doi.org/10.1136/bmj.n71>
21. Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *Journal of Strategic Information Systems*, 34(2), Article 101885. <https://doi.org/10.1016/j.jsis.2024.101885>
22. Prenio, J. (2024). Peering through the hype—Assessing suptech tools’ transition from experimentation to supervision (FSI Insights No. 58). Bank for International Settlements. <https://www.bis.org/fsi/publ/insights58.htm>



23. Prenio, J., Pustelnikov, A., & Yeo, J. (2024). Building a more diverse supotech ecosystem: Findings from surveys of financial authorities and supotech vendors (FSI Briefs No. 23). Bank for International Settlements. <https://www.bis.org/fsi/fsibriefs23.htm>
24. Reis, J. F., & Pinheiro Junior, L. P. (2025). Institutional Theory and Diffusion of Innovation: A theoretical approach to artificial intelligence. *BAR–Brazilian Administration Review*, 22(4), Article e250060. <https://doi.org/10.1590/1807-7692bar2025250060>
25. World Bank. (2025). Artificial intelligence for financial sector supervision: An emerging market and developing economies perspective. <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/099110525115015626>