

ARTIFICIAL INTELLIGENCE-ENABLED REGTECH FOR KYC, FINANCIAL COMPLIANCE, AND REGULATORY REPORTING: APPLICATIONS, RISKS, AND FUTURE DIRECTIONS

Yun Li*, HongZhe Zhang

School of Finance and Economics, Hainan Vocational University of Science and Technology, Haikou 571126, Hainan, China.

**Corresponding Author: Yun Li*

Abstract: Financial institutions face growing volumes of customer, payment, ownership, and regulatory data while remaining accountable for risk-based compliance decisions. Artificial intelligence-enabled regulatory technology can assist this work by combining digital identity services, entity resolution, rules, machine learning, graph analytics, natural language processing, and workflow automation. This narrative review explains the technical architecture of such systems and examines their use in know-your-customer procedures, financial compliance monitoring, investigations, and regulatory reporting. The technology can improve data reconciliation, prioritise alerts, reveal relationships that are difficult to detect in isolated records, retrieve evidence, and automate repeatable reporting steps. Its value, however, is constrained by incomplete labels, severe class imbalance, concept drift, adversarial adaptation, opaque models, privacy risks, fragmented rules, legacy systems, and dependence on external vendors. A lower false-positive rate does not by itself demonstrate improved monitoring effectiveness, and a suspicious transaction or activity report should be understood as an indicator for further review rather than definitive evidence of non-compliance. Responsible implementation therefore requires a layered architecture in which traceable controls preserve explicit regulatory requirements, models are validated against operational objectives, consequential decisions receive documented human oversight, and every data transformation, model version, explanation, override, and report remains auditable. Future development is likely to centre on privacy-preserving collaboration, temporal graph models, machine-readable regulation, and source-grounded generative assistants. AI-enabled RegTech should be treated as governed decision support, not as an autonomous substitute for legal interpretation, customer due diligence, investigation, or regulatory accountability.

Keywords: Regulatory technology; Know your customer; Financial compliance; Regulatory reporting; Artificial intelligence

1 INTRODUCTION

Financial compliance is both a legal obligation and a data-and-decision workflow. Institutions must identify customers and owners, screen names, monitor transactions, investigate unusual behaviour, document decisions, and submit reports under changing rules. In this article, RegTech denotes digital infrastructure and analytical tools used to operationalise compliance duties, from ongoing surveillance and evidence management to the preparation of mandated disclosures [1]. It reflects a movement from manual, document-centred processes toward proportionate, data-driven controls.

RegTech normally serves regulated institutions, whereas supervisory technology, or SupTech, serves authorities that collect data, assess risk, and conduct supervision. They exchange information and may use similar methods, but have different decision rights. A Financial Stability Board review identified data, cloud services, application programming interfaces, and analytics as enablers, with data quality, cyber risk, skills, governance, and transparency as constraints [2].

Research also shows that RegTech is not one product category. A systematic review identifies a nexus among regulatory interpretation, organisational processes, data, and technology rather than a single algorithmic solution [3]. For know-your-customer (KYC) and financial compliance work, the relevant technologies include digital identity, record linkage, rules engines, supervised and unsupervised learning, graph analysis, and natural language processing. The Financial Action Task Force (FATF) recognises opportunities for new technologies in customer due diligence, monitoring, risk assessment, and communication with supervisors, but conditions their responsible use on governance, data protection, explainability, and human oversight [4].

FATF uses customer due diligence (CDD) as its standards term, whereas KYC is a broader industry label. Its standards address CDD, suspicious transaction reports (STRs), and confidentiality, but national implementation varies. This paper uses STR generally, recognises that some jurisdictions use suspicious activity report (SAR), and separates AML intelligence reports from prudential, statistical, or transaction-data submissions.

This selective narrative review examines technical principles, applications, problems, and future directions. Targeted searches of Crossref, publisher records, and official sites ended on 26 August 2026. Terms combined RegTech, KYC or CDD, AML, transaction monitoring, graph learning, explainability, privacy, and digital regulatory reporting. English sources were selected for relevance, verifiable metadata, primary authority, and clear evidence boundaries, but the

search was not exhaustive. The central argument is that AI creates value through prioritisation, evidence assembly, and repeatable execution inside a controlled workflow, without replacing institutional responsibility.

2 TECHNICAL ARCHITECTURE OF AI-ENABLED REGTECH

2.1 Data, Identity, and Entity Resolution

The foundation is a governed data layer. Inputs include identity, ownership, transaction, screening, external-intelligence, case, and regulatory records. Formats, identifiers, and update cycles differ. Before modelling, the system must validate fields, standardise values, preserve provenance, resolve duplicates, and record transformations.

Digital identity can support remote onboarding when its assurance level, governance, and safeguards are appropriate to the risk. FATF guidance does not treat non-face-to-face identification as automatically high risk or endorse a particular technology [5]. Institutions must still assess the reliability of evidence, the accuracy of identity verification, the integrity of transactions, inclusiveness, privacy protection, and cybersecurity.

Entity resolution links imperfect records that may refer to the same person or organisation. Many cases require probabilistic comparison of names, addresses, dates, corporate roles, and network relationships. Output should expose candidate matches, contributing fields, uncertainty, and an escalation path. A mistaken merge can contaminate a risk profile, while a missed link can conceal ownership or transaction relationships.

2.2 Hybrid Rules and Machine Learning

Most practical architectures are hybrid. Deterministic rules encode legal thresholds and mandatory controls that officers must trace to policy. Models can then rank alerts, estimate risk from more variables, or identify unanticipated patterns. A review describes functions including typology recognition, link and behavioural analysis, risk scoring, anomaly detection, and geographic analysis.

Supervised models learn from historical outcomes such as investigator dispositions or STRs. They may improve queue ordering, but these labels are incomplete and institution-specific. A reported customer does not necessarily indicate confirmed non-compliance, while an unreported case does not necessarily indicate full compliance. In one bank-specific study, deep sequence models were tested as second-stage qualifiers of existing alerts and, in a smaller experiment, as generators of new alerts. These results cannot be assumed to transfer to another bank, period, or jurisdiction.

Unsupervised and semi-supervised methods are useful when reliable positive labels are scarce. Clustering, autoencoders, and related methods learn common behaviour and surface unusual observations for review. An anomaly score identifies deviations from established patterns rather than directly indicating compliance status, and therefore requires appropriate business context for interpretation. A rare transaction may be fully legitimate, while recurring patterns of non-compliant activity may appear statistically ordinary when they are well represented in the training data.

2.3 Graph Analytics, Language Processing, and Human Review

Money movement is relational. Graphs represent customers, accounts, companies, directors, devices, addresses, and external counterparties as nodes, with transactions and other relationships as edges. Traditional social-network measures can identify intermediaries, dense clusters, shared ownership, or unusual cross-industry connections that account-level features miss. More recent heterogeneous graph neural networks can learn from several node and edge types at once. A large single-bank study showed that this approach can model transaction and business-role relationships, although unobserved external links and incomplete labels remain important limitations [6].

Table 1 Layered Architecture for AI-Enabled RegTech

Layer	Main inputs	AI or automation function	Operational output	Essential control
Data and identity	KYC records, identifiers, ownership, transactions, external lists	Validation, standardisation, document extraction, entity resolution	Governed customer and relationship view	Source lineage, time stamps, match uncertainty, access control
Detection	Rules, customer profiles, transaction sequences	Risk scoring, alert ranking, anomaly detection	Prioritised alerts and risk indicators	Rule traceability, time-based validation, threshold approval
Network and text	Graph relationships, adverse media, regulations, case notes	Community detection, graph learning, entity and obligation extraction	Related parties, patterns, evidence candidates	Provenance, relevance checks, no inference from text alone
Investigation	Alerts, documents, explanations, prior decisions	Evidence retrieval, chronology building, draft narrative support	Reviewable case package	Investigator sign-off, override reasons, quality assurance
Reporting	Approved facts, reporting rules, taxonomies, templates	Mapping, validation, calculation, form population	Draft regulatory submission	Reconciliation, version control, legal and management approval

Natural language processing (NLP) can extract names and obligations, classify adverse news, link entities, and search case files. Watch-list research has combined string similarity with KYC and transaction attributes to reclassify alerts. This supports assisted review, not autonomous adverse action.

These components should converge in a case-management layer. The investigator needs the triggering rule or model score, relevant transactions, network context, source documents, uncertainty, prior actions, and an explanation of the features that influenced prioritisation. The interface should support requests for further evidence, documented overrides, quality review, and escalation. Table 1 summarises the architecture.

3 APPLICATION SCENARIOS

3.1 Digital Onboarding and Continuous KYC

At onboarding, automation can extract document fields, check completeness, perform lawful identity comparisons, screen names, and assemble ownership information. Entity resolution can reduce duplicates and focus analysts on ambiguous matches. A risk engine can recommend a due-diligence level from customer, geography, product, ownership, activity, and screening factors.

The recommendation should initiate questions rather than close the case. Analysts need to review weak evidence, unexplained ownership, conflicting records, and accessibility barriers. Continuous KYC can monitor ownership changes, document expiry, adverse media, or activity inconsistent with the relationship's stated purpose. Event-driven review can be timelier than a periodic calendar, but noisy data can create perpetual low-value alerts.

3.2 Transaction Monitoring and Investigation

Rules remain useful for known typologies and explicit policy conditions, while models can order alerts according to likely relevance or search for new patterns. The operational objective is not simply the highest classification accuracy. Investigators work with limited capacity, extreme class imbalance, and costly errors in both directions. Takáts's analysis of "crying wolf" explains how stronger reporting incentives can produce excessive reports, dilute their information value, and reduce enforcement effectiveness.

Industry interviews identify the need to improve monitoring efficiency, adapt to evolving transaction patterns, enhance explainability, strengthen scalability, and increase flexibility as recurring requirements for advanced transaction monitoring systems. A real-bank study of automated false-positive suppression likewise found a trade-off: stronger suppression can remove useful alerts as well as unproductive ones. Precision, recall, and calibration measured against investigator dispositions or regulatory filings primarily reflect the consistency of model outputs with reference labels rather than directly measuring monitoring effectiveness. These metrics should therefore be interpreted alongside alert volume, investigation time, escalation yield, sampled reviews of cases receiving different monitoring scores, and delayed official outcomes where available. A lower alert volume should be understood as a workflow outcome and does not, by itself, demonstrate improved monitoring effectiveness.

AI can assemble transactions, visualise networks, retrieve prior cases, and draft a chronology or STR narrative. Generated text may omit qualifying evidence or invent connections. The draft must cite case records and distinguish observation from inference. A filing decision must follow applicable law and the institution's accountable process, and should not be delegated to an unconstrained language model.

3.3 Regulatory Change and Reporting Automation

Regulatory reporting involves a chain from legal text to interpretation, data definitions, calculations, validation rules, forms, and approval. AI and semantic technologies can help identify changed provisions, map concepts, locate affected controls, and route tasks to subject-matter experts. Butler and O'Brien's RegTech case study shows how semantic standards can connect regulatory concepts, reporting rules, and enterprise data, while also illustrating that expert interpretation remains central [7].

For stable and precisely defined obligations, machine-readable and executable instructions can automate more of the chain. A joint Financial Conduct Authority and Bank of England pilot explored standardised firm data, executable reporting instructions, and automatic report generation in two limited use cases [8]. It demonstrated prototype feasibility, not a production-ready or universally applicable system. Ambiguous legal concepts, exceptions, and jurisdiction-specific judgments cannot simply be converted into code without governance.

The European Banking Authority's market analysis identified financial compliance, financial integrity, prudential reporting, information-security governance, and creditworthiness assessment among the most common areas of RegTech application [9]. Institutions and providers reported benefits such as improved monitoring and fewer manual errors, but the evidence was survey-based. Reporting automation should therefore be evaluated through reconciled data quality, correction rates, timeliness, change-implementation effort, and audit findings rather than vendor claims alone.

3.4 RegTech and SupTech Interaction

Authorities increasingly receive granular data and use validation, network analysis, NLP, and risk dashboards to select supervisory priorities. This is SupTech even when the method resembles a bank's RegTech system. A sound interface

can improve feedback on data errors and reduce repeated information requests. It should not blur accountability: the institution remains responsible for complete and accurate submissions, and the authority remains responsible for supervisory judgment. Common definitions, secure interfaces, and transparent validation rules are more valuable than forcing both sides to use the same proprietary model.

4 RISKS AND CHALLENGES

4.1 Data Quality, Labels, and Performance Measurement

Models inherit incomplete records, inconsistent identifiers, delayed dispositions, and historical policy choices. Reference outcomes are rarely fully available, while investigator labels may also reflect operational priorities and the characteristics of previously adopted models. Random splits may leak related customers or later information. Validation should use entity-level separation, out-of-time testing, scenarios, and review of low-ranked cases.

Class imbalance makes accuracy misleading. A model can classify almost every observation as normal and still appear accurate. Label-relative metrics should match the use case, while sampled review of non-alerted and low-scored cases probes what proxy labels miss. Thresholds allocate investigative capacity and risk, and the trade-off observed in alert-suppression research makes them governance choices rather than merely technical settings.

4.2 Concept Drift and Adversarial Adaptation

Customer behaviour, payment channels, regulatory requirements, and evolving transaction patterns continue to change, highlighting the need for adaptive approaches, a consideration also emphasized by financial compliance practitioners. Data drift concerns inputs, concept drift the relationship between inputs and outcomes, and policy drift the institution's treatment of risk. Monitoring should compare scores, feature quality, alert composition, investigator outcomes, and subgroup errors over time. Material change requires diagnosis and controlled recalibration.

Adversaries may split transactions, recruit intermediaries, imitate normal behaviour, manipulate identity evidence, or probe thresholds. Publicly exposing detailed decision logic can therefore create gaming risk, but secrecy cannot replace internal explainability and independent validation. Robustness testing should include altered identities, new networks, missing data, coordinated low-value activity, and attacks on document or language-processing components.

4.3 Explainability, Fairness, and Accountability

An explanation must be useful for its audience. Investigators need actionable case factors and supporting records, validators need model logic and sensitivity evidence, and supervisors need governance and audit trails. Customer-facing decisions may require an understandable reason where disclosure is lawful, but STR confidentiality and anti-tipping-off rules constrain disclosure. Explainability techniques can be inaccurate, unstable, or misleading and should not be treated as causal accounts [10]. Where performance is comparable, interpretable models may be preferable for high-impact decisions. More complex models require stronger validation, documentation, and human review.

Because KYC controls affect access to financial services, institutions should test error patterns across legally appropriate groups, examine accessibility and identity-match failures, and assess whether variables are relevant and proportionate to risk [4]. Removing a protected attribute does not remove proxy effects. Fairness criteria can conflict, so the chosen standard needs legal and policy justification rather than an automated optimisation target.

Accountability cannot be outsourced. The Financial Stability Board identifies model risk, data quality, governance, cyber risk, and dependence on concentrated third-party providers as possible financial-sector vulnerabilities from AI [11]. A bank that buys a model remains responsible for knowing its permitted use, training-data constraints, update process, performance, failure modes, and exit options.

4.4 Privacy, Security, and Fragmented Requirements

KYC and AML systems combine identity, financial, relationship, and sometimes biometric or adverse-media data. Centralising these sources may improve detection while increasing the impact of unauthorised access or erroneous linkage [2]. Controls should include purpose limitation, data minimisation, encryption, role-based access, retention schedules, secure testing environments, and records of downstream use. Confidential customer or case data should not enter an external generative service without approved contractual, technical, and jurisdictional safeguards [11].

Fragmented definitions and reporting formats add another constraint. The same concept may have different thresholds, exceptions, time windows, or evidence requirements across jurisdictions. Legacy systems may encode local meanings that are poorly documented. Premature automation can reproduce one mistaken interpretation at scale. A shared semantic layer helps, but every mapping still needs ownership, versioning, effective dates, and a route for resolving ambiguity [3].

5 RESPONSIBLE IMPLEMENTATION AND GOVERNANCE

A responsible programme begins with the compliance objective, not the model. Each use case should identify the governing obligation, affected customers, decision owner, acceptable automation, required evidence, and consequence

of error. Human involvement should be calibrated to applicable law, uncertainty, and consequence. Low-impact extraction or reconciliation can be highly automated, while consequential actions need documented oversight and escalation or approval points where required. The institution remains accountable, and STR handling must respect confidentiality and anti-tipping-off requirements.

Data and model governance should cover lineage, lawful use, quality thresholds, feature definitions, training and validation periods, version history, access, and retention. Independent validation should reproduce the data pipeline, test temporal and subgroup performance, challenge assumptions, inspect low-scored cases, and compare the model with rules and simpler alternatives. Performance monitoring needs pre-agreed triggers for investigation, rollback, recalibration, or retirement.

Operational design is equally important. Investigators should see why a case was prioritised, what evidence is missing, and how to override a recommendation. Overrides should be encouraged when justified and analysed for recurring model gaps. Quality assurance should sample both escalated and non-escalated cases. Regulatory reporting needs source reconciliation, maker-checker controls, change testing, and tamper-evident, version-controlled records that preserve corrections and history. Table 2 links lifecycle stages to controls and evidence.

Table 2 Governance Controls Across the RegTech Lifecycle

Stage	Key question	Control	Evidence retained
Use-case approval	Is automation appropriate for the obligation and consequence?	Risk classification, legal review, named decision owner	Scope, prohibited uses, approval record
Data preparation	Are inputs lawful, accurate, current, and traceable?	Quality rules, lineage, access and retention controls	Source, transformations, exceptions, data version
Development and validation	Does the method improve the defined workflow without hiding risk?	Baseline comparison, out-of-time and subgroup tests, stress tests	Code or specification, metrics, limitations, validation findings
Deployment	Can users understand, contest, and safely override output?	Human decision points, role-based interface, fallback process	Model version, recommendation, explanation, user action
Monitoring	Has data, performance, policy, or behaviour changed?	Drift indicators, threshold review, sampled review of low-scored and non-alerted cases	Trends, incidents, overrides, remediation decisions
Change and retirement	Can the institution reproduce and reverse a change?	Controlled release, rollback, vendor exit and record retention	Test results, approvals, release log, archived artefacts

Governance metrics should resist a narrow return-on-investment view. Useful measures include corrected data errors, time to obtain missing evidence, label-relative alert recall at fixed capacity, sampled low-score findings, report corrections, overdue reviews, model incidents, and decisions with complete provenance. Efficiency matters, but only alongside effectiveness, customer impact, and defensibility.

6 FUTURE DIRECTIONS

Cross-institution collaboration is promising because one institution sees only part of a cross-border network. Centralised data pooling broadens visibility but also concentrates sensitive data. Distributed analytics and privacy-enhancing methods such as secure multiparty computation and homomorphic encryption may reduce some exposure [12]. None makes data use automatically lawful or secure. Governance must address purpose, access, inconsistent inputs, leakage and re-identification tests, responsibility, and redress.

Building on heterogeneous graph evidence, future work could model the timing as well as the types of relationships among people, companies, accounts, and devices [6]. Evaluation should cover new entities, delayed outcomes, missing external edges, and investigator usefulness. Hybrid designs remain suitable: rules protect mandatory coverage, models find candidates, and analysts test explanations against evidence.

Machine-readable regulation and granular reporting may reduce repeated manual interpretation when legal concepts, data definitions, and validation rules are linked through shared standards [7]. Progress depends less on a universal AI translator than on maintained taxonomies, versioned rule representations, common identifiers, and formal processes for ambiguity. Regulators and institutions can then test changes in controlled environments before they affect submissions. Generative AI may become a compliance copilot for searching approved sources, comparing rule versions, summarising a case, or drafting a report narrative. This paper proposes retrieval-grounded generation that displays document identity, effective date, passage, and data lineage. Calculation engines and rule services should execute formal logic, while the language model explains or assembles rather than invents results. High-impact outputs require review, and system prompts are not a substitute for access control, validation, or audit logs [11].

Finally, evaluation needs more comparable evidence. Institutions and researchers should document datasets, label construction, time periods, baselines, operational capacity, subgroup performance, and human-review outcomes. Privacy-assessed benchmark data and controlled sandboxes can improve reproducibility, but synthetic data and proofs of concept still require validation on current, institution-specific data before deployment.

7 REVIEW LIMITATIONS

This targeted narrative synthesis does not claim exhaustive study identification or quantitative aggregation. Public evidence is uneven because customer and investigation data are confidential, institutions use different labels and thresholds, and operational failures are rarely disclosed in comparable form. Several influential sources are policy reports, surveys, case studies, prototypes, or single-institution analyses. Their findings support the technical feasibility and governance requirements of these approaches, while indicating that their accuracy, cost efficiency, and broader effectiveness may vary across different contexts. Rapid changes in models, payment systems, and regulation also limit the durability of specific implementations.

8 CONCLUSION

AI-enabled RegTech can strengthen KYC, AML, and regulatory reporting when it connects governed data with appropriate rules, statistical models, network analysis, language processing, and accountable workflows. It can reconcile records, prioritise alerts, reveal relationships, assemble evidence, and automate stable reporting steps. None of these functions transfers the institution's responsibility for compliance outcomes.

The strongest implementation pattern is layered and risk-based. Deterministic controls preserve traceable obligations, models support prioritisation and pattern discovery, and human oversight is calibrated to applicable law, uncertainty, risk, and consequence. Validation must address incomplete labels, class imbalance, time change, subgroup effects, adversarial behaviour, and operational capacity. Data lineage, model versions, explanations, overrides, and reporting transformations must remain auditable. Privacy, cybersecurity, interoperability, and vendor dependence are design constraints rather than issues to address after deployment.

Future systems may support privacy-preserving collaboration, richer temporal graphs, executable reporting instructions, and source-grounded generative assistance. Their success should be judged by defensible risk coverage, reliable evidence, appropriate customer treatment, and accountable decisions, not by automation rates alone. RegTech is therefore best understood as governed compliance infrastructure that augments professional judgment rather than an autonomous decision maker.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

FUNDING

This work was supported in part by the research project "Research on the Mechanism of Digital Empowerment for Developing the Business Environment of the Free Trade Port" (No. 2026-ZCKT-190).

REFERENCES

- [1] Arner D W, Barberis J, Buckley R P. FinTech, RegTech, and the reconceptualization of financial regulation. *Northwestern Journal of International Law & Business*, 2017, 37(3): 371-414. <https://scholarlycommons.law.northwestern.edu/njilb/vol37/iss3/2/>.
- [2] Financial Stability Board. *The Use of Supervisory and Regulatory Technology by Authorities and Regulated Institutions: Market Developments and Financial Stability Implications*. Basel: Financial Stability Board, 2020. <https://www.fsb.org/uploads/P091020.pdf>.
- [3] Grassi L, Lanfranchi D. RegTech in public and private sectors: The nexus between data, technology and regulation. *Journal of Industrial and Business Economics*, 2022, 49(3): 441-479. DOI: 10.1007/s40812-022-00226-0.
- [4] Financial Action Task Force. *Opportunities and Challenges of New Technologies for AML/CFT*. Paris: Financial Action Task Force, 2021. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Opportunities-Challenges-of-New-Technologies-for-AML-CFT.pdf.coredownload.pdf>.
- [5] Financial Action Task Force. *Guidance on Digital Identity*. Paris: Financial Action Task Force, 2020. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Guidance-on-Digital-Identity.pdf.coredownload.pdf>.
- [6] Johannessen F, Jullum M. Finding money launderers using heterogeneous graph neural networks. *The Journal of Finance and Data Science*, 2025, 11: 100175. DOI: 10.1016/j.jfds.2025.100175.
- [7] Butler T, O'Brien L. Understanding RegTech for digital regulatory compliance. In: Lynn T, Mooney J G, Rosati P, et al., editors. *Disrupting Finance: FinTech and Strategy in the 21st Century*. Cham: Palgrave Macmillan, 2019: 85-102. DOI: 10.1007/978-3-030-02330-0_6.
- [8] Financial Conduct Authority, Bank of England. *Digital Regulatory Reporting: Pilot Phase 1 Report*. London: Financial Conduct Authority and Bank of England, 2018. <https://www.fca.org.uk/publication/discussion/digital-regulatory-reporting-pilot-phase-1-report.pdf>.
- [9] European Banking Authority. *EBA Analysis of RegTech in the EU Financial Sector*. EBA/REP/2021/17. Paris: European Banking Authority, 2021.

- <https://www.eba.europa.eu/publications-and-media/press-releases/eba-assesses-benefits-challenges-and-risks-regtech-use-eu-and>.
- [10] Perez-Cruz F, Prezio J, Restoy F, et al. Managing explanations: how regulators can address AI explainability. FSI Occasional Papers No 24. Basel: Bank for International Settlements, 2025. <https://www.bis.org/fsi/fsipapers24.htm>.
- [11] Financial Stability Board. The Financial Stability Implications of Artificial Intelligence. Basel: Financial Stability Board, 2024. <https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/>.
- [12] Financial Action Task Force. Stocktake on Data Pooling, Collaborative Analytics and Data Protection. Paris: Financial Action Task Force, 2021. <https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Stocktake-Datapooling-Collaborative-Analytics.pdf>.