



Artificial Intelligence-Assisted Anti-Money Laundering Compliance for Small and Medium-Sized Businesses: A Risk-Based Governance Framework

S. M. Sadek Ismail

St. John's University Law School

sadek.ismail219@gmail.com

+1 347-258-2297

DOI: <https://doi.org/10.67487/ijfm.v1i1.289>

Citation: Ismail S. M. S. (2026). Artificial Intelligence-Assisted Anti-Money Laundering Compliance for Small and Medium-Sized Businesses: A Risk-Based Governance Framework. *International Journal of Finance, Management and Governance*, 2(3) <https://doi.org/10.67487/ijfm.v1i1.289>

Abstract

Artificial intelligence (AI) is increasingly presented as a solution to persistent weaknesses in anti-money-laundering (AML) compliance, including high false-positive rates, fragmented customer data, delayed investigations, and limited compliance capacity. Yet most discussions focus on banks and large financial institutions, leaving small and medium-sized businesses (SMEs), fintech firms, money-service businesses, professional service providers, and other regulated or exposed enterprises without a proportionate implementation model. This conceptual article examines how AI can support business-facing AML functions while preserving legal accountability, explainability, privacy, and human judgment. The article identifies five high-value use cases: customer due diligence, transaction monitoring, adverse-media screening, network analysis, and investigation support. It does this by using international AML standards, regulatory guidance, AI-risk governance frameworks, and recent financial-crime technology literature. It also evaluates major risks, including biased data, opaque models, automation bias, model drift, cyber manipulation, privacy violations, and the use of generative AI by criminals. The article proposes the Responsible AI-AML Governance Framework, organized around six pillars: risk scoping, data governance, model validation, human oversight, documentation and auditability, and continuous monitoring. The framework is designed as a proportionate roadmap for businesses with limited resources. The article concludes that AI should augment, rather than replace, accountable compliance professionals and



that effective adoption depends more on governance quality than algorithmic sophistication.

Keywords: artificial intelligence; anti-money laundering; compliance; small and medium-sized enterprises; governance; financial crime

1. Introduction

Money laundering remains a central threat to the integrity of financial markets, commercial systems, and public institutions. Criminal proceeds are increasingly moved through complex combinations of bank accounts, shell entities, trade transactions, digital assets, payment platforms, professional intermediaries, and ordinary businesses. At the same time, firms subject to anti-money-laundering and counter-terrorist-financing obligations face growing volumes of customer information, transaction data, sanctions updates, adverse media, and regulatory expectations. Traditional rule-based compliance systems frequently generate large numbers of alerts that require manual review, even though only a small proportion may ultimately justify escalation. This creates a difficult operational environment in which compliance teams must identify rare but consequential risks within large, noisy datasets.

Artificial intelligence has emerged as a potential response to this problem. Machine-learning models can identify non-linear patterns, detect unusual behavior, prioritize alerts, link apparently unrelated entities, and extract information from unstructured documents. Natural-language-processing tools can assist with adverse-media review, customer-risk narratives, policy analysis, and case summarization. Graph analytics can reveal networks of counterparties and beneficial owners that are difficult to detect through isolated transaction rules. These capabilities are especially attractive to businesses seeking to improve effectiveness while controlling the cost of compliance.

The promise of AI, however, should not be confused with automatic legal compliance. An algorithm may amplify errors in training data, produce outputs that cannot be explained to regulators, overlook new criminal typologies, or unfairly burden customers from particular locations or demographic groups. Generative AI also creates new threats by enabling synthetic identities, deepfake documents, scalable fraud communications, and automated evasion strategies. Consequently, the same technological development can strengthen both financial-crime detection and financial-crime execution.

A further gap concerns organizational size. Regulatory and academic discussions often assume the availability of large datasets, specialist model-risk teams, advanced technology infrastructure, and substantial compliance budgets. Those assumptions do not fit many SMEs,



fintech start-ups, money-service businesses, accountants, lawyers, real-estate businesses, payment agents, and cross-border traders. Smaller organizations may be exposed to money-laundering risk but lack the resources to design, validate, and monitor complex AI systems. A governance approach that merely copies the practices of global banks may therefore be impractical and may create a false appearance of control.

This article develops a proportionate framework for AI-assisted AML compliance in business settings. It addresses three research questions: (1) Which AML functions are most suitable for AI assistance? (2) What legal, ethical, operational, and governance risks arise when businesses use AI for AML purposes? and (3) How can SMEs and other resource-constrained organizations implement AI responsibly without weakening human accountability? The article contributes by integrating AML risk management with recognized AI-governance principles and translating them into a practical six-pillar model.

2. Literature Review and Theoretical Foundation

The modern AML system is based on a risk-based approach. Rather than treating every customer and transaction identically, organizations are expected to identify, assess, and understand their exposure and then apply proportionate controls. The Financial Action Task Force (FATF) has emphasized that new technologies can improve the speed, quality, and efficiency of customer due diligence, transaction monitoring, risk assessment, and supervisory communication. This position supports innovation, but it also requires technology to be consistent with privacy, cybersecurity, and broader legal obligations.

Traditional AML transaction-monitoring systems generally rely on scenarios and thresholds. Examples include repeated cash deposits below reporting limits, rapid movement of funds through newly opened accounts, transactions involving high-risk jurisdictions, or unusual changes in customer activity. These rules are understandable and auditable, but they are often static. Criminal behavior can adapt around thresholds, while legitimate customers may trigger alerts because their circumstances differ from the assumptions embedded in a rule. Machine learning can supplement these systems by learning patterns from historical cases or by identifying anomalies without requiring every pattern to be specified in advance.

Supervised learning uses labeled historical examples, such as alerts previously classified as suspicious or non-suspicious. It may improve alert prioritization when labels are accurate and sufficiently representative. Unsupervised learning looks for clusters, outliers, or unusual relationships without relying entirely on prior labels. This is useful when emerging typologies differ from known cases, although anomalies are not necessarily illegal. Semi-supervised methods combine a smaller amount of labeled data with a larger amount of unlabeled data.



Graph-based methods analyze relationships among customers, accounts, devices, addresses, companies, and counterparties, reflecting the networked nature of laundering schemes.

Natural-language processing expands the scope of analysis beyond numerical transactions. It can extract names, entities, locations, allegations, and risk indicators from corporate records, news reports, customer documents, internal case notes, and regulatory materials. Large language models may summarize lengthy files and draft investigation narratives. Nevertheless, generative systems can hallucinate facts, omit qualifications, or produce confident but unsupported conclusions. Their output must therefore be treated as an assistive draft rather than a verified compliance determination.

The theoretical foundation of responsible AI-AML adoption combines the risk-based approach with socio-technical governance. A socio-technical system includes not only software but also data sources, human reviewers, incentives, procedures, vendors, management oversight, and the legal environment. Socio-technical systems theory recognizes that organizational performance depends on the interaction between social and technical elements (Trist & Bamforth, 1951; Baxter & Sommerville, 2011). Failures may arise even when a model is technically accurate—for example, when staff misunderstand a risk score, when escalation procedures are weak, or when business pressure discourages investigation. Effective governance must therefore evaluate the entire decision process.

The NIST AI Risk Management Framework offers a useful organizing logic through its functions of Govern, Map, Measure, and Manage. It emphasizes validity, reliability, safety, resilience, transparency, explainability, privacy, and fairness. These qualities align closely with AML expectations concerning internal controls, independent testing, recordkeeping, risk assessment, and accountable decision-making. However, a general AI framework does not by itself specify how to conduct customer due diligence, triage alerts, document suspicious-activity decisions, or manage vendor AML tools. The framework proposed in this article adapts those principles to the AML context.

3. Methodology

This study uses an integrative conceptual review and framework-development methodology. It does not claim to report primary empirical findings or the performance of a proprietary model. The analysis synthesizes four categories of materials: international AML standards and policy reports; official regulatory and supervisory guidance; recognized AI-risk-management frameworks; and peer-reviewed or institutional research on machine learning, anomaly detection, network analytics, and financial-crime compliance.



The materials were examined to identify recurring themes concerning technological capability, implementation barriers, legal accountability, and organizational governance. The analysis then mapped AI use cases to stages of the AML compliance lifecycle: onboarding, customer-risk assessment, ongoing monitoring, alert investigation, reporting, and quality assurance. Risks were categorized as data risks, model risks, human-factor risks, legal and ethical risks, operational risks, and adversarial risks. Finally, the categories were consolidated into a six-pillar governance framework designed to be adaptable to different business sizes and risk profiles.

The conceptual method is appropriate because many organizations cannot publicly disclose detailed AML models, suspicious-activity cases, or law-enforcement outcomes. It also avoids presenting unverified numerical claims about cost savings or detection accuracy. The limitation is that the proposed framework requires future empirical testing across industries and jurisdictions. Its purpose is to provide a theoretically grounded implementation model rather than a universal technical specification.

4. AI Use Cases in Business AML Compliance

4.1 Customer due diligence and onboarding: AI can support identity verification, document classification, extraction of corporate ownership data, duplicate detection, and comparison of customer information across sources. Optical and language models may identify inconsistencies between an application, identification document, corporate registry entry, and declared business activity. Risk models can combine factors such as geography, product type, ownership complexity, expected transaction behavior, delivery channel, and negative information. For smaller businesses, the greatest benefit may come from structured workflow tools that ensure required information is collected and higher-risk cases receive enhanced review.

AI should not convert customer risk into an unexplained score. A reviewer must be able to identify the principal factors that increased or reduced risk. The business should distinguish objective legal triggers, such as sanction matches, from probabilistic indicators, such as behavioral similarity to prior cases. Customers should not be rejected solely because a model identifies a weak statistical association. Where the consequences are significant, an appropriately trained employee should review the underlying evidence and document the decision.

4.2 Transaction monitoring and alert prioritization: Machine learning can supplement rule-based monitoring by identifying unusual changes relative to a customer's own history and comparable customer groups. It can also combine multiple weak indicators that would not independently cross a threshold. A practical model may rank alerts according to risk rather than automatically closing them. This can direct scarce investigative resources toward cases with stronger combinations of indicators while preserving rules required by law or policy.

A central danger is label quality. Historical suspicious-activity reports are not equivalent to proven money laundering, and past investigator decisions may contain inconsistent judgments. Training a model on those labels can reproduce old errors. Moreover, legitimate activity evolves. A model trained before a major economic disruption, migration pattern, product launch, or payment innovation may treat new normal behavior as suspicious. Ongoing monitoring for drift is therefore necessary.

4.3 Adverse-media and sanctions screening: Natural-language tools can search large volumes of multilingual news and identify possible connections to fraud, corruption, sanctions evasion, trafficking, cybercrime, or other predicate offenses. Entity-resolution tools can reduce false matches by considering date of birth, location, occupation, and associated organizations. A human reviewer must still assess source reliability, distinguish allegations from convictions, and avoid confusing individuals with similar names.

4.4 Network and beneficial-ownership analysis: Laundering commonly involves chains of entities, common directors, shared addresses, nominee owners, devices, or transaction counterparties. Graph analytics can reveal communities and central actors that are not apparent from account-by-account review. This is particularly valuable for trade businesses, payment companies, and professional service providers that encounter networks of related parties. Even a modest organization can use graph visualization to organize known relationships, provided that the data are lawfully obtained and accurately maintained.

4.5 Investigation and reporting support. Generative AI can summarize customer files, prepare chronologies, identify missing information, and draft a first version of an internal case narrative. It may also help compliance personnel compare observed facts with known typologies. The final decision to escalate, restrict, terminate, or report activity must remain with authorized personnel. Confidential customer or suspicious-activity information should not be entered into public generative-AI services unless the organization has established lawful, secure, and contractually controlled use.

5. Risks and Governance Challenges

Data quality is the first and often most important risk. AI models cannot reliably correct incomplete customer records, inconsistent transaction codes, duplicated identities, outdated beneficial-ownership information, or poorly documented investigations. A sophisticated model trained on weak data may produce an appearance of precision without genuine reliability. Organizations should therefore treat data remediation as part of AML control design, not as a preliminary technical task that can be delegated entirely to a vendor.

Bias and disparate impact also require attention. Geographic location, language, nationality-related information, occupation, or transaction patterns may act as proxies for protected or vulnerable characteristics. Some variables are legitimately relevant to AML risk, but their use must be connected to documented risk evidence rather than stereotypes. A business should test whether the model disproportionately escalates particular groups and determine whether differences are justified by valid risk factors. Fairness does not require identical outcomes; it requires defensible, evidence-based treatment and meaningful review.

Explainability is essential because compliance decisions must be communicated to investigators, senior management, auditors, regulators, and sometimes customers. Highly complex models may improve predictive performance but make individual decisions difficult to interpret. A proportionate system should provide reason codes, relevant transactions, key relationships, and supporting evidence. Where a model cannot explain why a case is high risk, it should not serve as the sole basis for consequential action.

Automation bias occurs when employees defer to a system even when evidence suggests the output is wrong. The opposite problem, automation distrust, occurs when investigators ignore a system because it produces too many weak alerts or because they do not understand it. Training, user-interface design, feedback mechanisms, and performance reporting influence whether human oversight is real or merely formal. Reviewers must have authority to disagree with the model, and disagreements should be captured for quality improvement.

Model drift and typology change are unavoidable. Criminals adapt to controls, customers change behavior, products evolve, and data pipelines break. Performance at deployment does not guarantee continuing effectiveness. Monitoring should include changes in alert volume, override rates, false-positive indicators, missed-case reviews, data completeness, and performance across customer segments. Material changes should trigger revalidation.

Vendor dependency creates additional risk for SMEs. A vendor may describe a product as AI-powered without disclosing training data, validation methods, limitations, or subcontractors. The purchasing business remains accountable for how the tool is used. Contracts should address data ownership, confidentiality, security incidents, model changes,

audit rights, service continuity, retention, deletion, and regulatory cooperation. Businesses should avoid relying on marketing claims as evidence of effectiveness.

Adversarial use of AI is a growing concern. Criminals may generate synthetic identification, manipulate invoices, create deepfake video or audio, automate social engineering, and test transactions against perceived monitoring thresholds. FinCEN has warned financial institutions about deepfake-enabled fraud, while FATF has highlighted AI-related financial-crime vulnerabilities. Defensive systems must therefore include document-authenticity checks, device intelligence, out-of-band verification, staff awareness, and escalation for synthetic-media indicators.

Privacy and confidentiality constrain the collection and reuse of data. AML obligations may authorize or require certain processing, but organizations should still apply purpose limitation, access controls, retention rules, encryption, and secure deletion. Cross-border transfers and cloud services may create additional legal requirements. Data pooling among institutions can improve detection of networks, but it must be structured within applicable privacy, secrecy, and information-sharing laws.

Finally, accountability cannot be outsourced to a model. Senior management must understand the role of AI in the compliance program, approve the risk appetite, allocate resources, and receive meaningful reporting. Compliance officers need sufficient independence to challenge business objectives that create pressure to reduce alerts or accelerate onboarding. Independent testing should evaluate both technical performance and whether the system supports the organization's legal obligations.

6. The Responsible AI-AML Governance Framework

The proposed Responsible AI-AML Governance Framework contains six pillars. It is designed as a scalable control structure rather than a single technology architecture.

Pillar 1: Risk scoping and proportionality. The organization should begin with a documented AML risk assessment and a clearly defined problem. It should identify the customers, products, jurisdictions, channels, and transaction types within scope; the decision that the AI tool will support; and the potential harm from false positives and false negatives. AI should be selected only where it offers a material improvement over simpler rules or workflow controls. A small business may appropriately use AI for document extraction and alert ranking without deploying autonomous customer rejection.

Pillar 2: Data governance. The business should maintain a data inventory showing source, ownership, lawful basis, quality, retention, access, and permitted use. Critical data fields should have validation rules. Training, testing, and production data should be separated where appropriate. Sensitive information should be minimized and protected. Data-quality issues must be tracked for remediation rather than silently absorbed by the model.

Pillar 3: Model development, validation, and change control. Before use, the organization should document the model's purpose, logic, inputs, outputs, assumptions, limitations, and performance measures. Validation should be independent from the developer to the extent proportionate to the organization's size. For vendor tools, independent validation may involve sample testing, benchmarking, reason-code review, and contractual access to documentation. Changes to thresholds, features, or model versions should be approved and recorded.

Pillar 4: Human oversight and decision rights. Policies should state which decisions require human approval and who has authority to override the model. High-impact outcomes, such as filing a suspicious-activity report, terminating a relationship, freezing activity where legally permitted, or rejecting a customer, should not be made solely by an opaque model. Reviewers should receive training on model strengths, limitations, and common error patterns.

Pillar 5: Documentation, explainability, and auditability. Each material output should be traceable to the data and model version used. The system should preserve reason codes, relevant evidence, reviewer actions, overrides, and escalation outcomes. Documentation should be understandable to compliance professionals rather than only data scientists. Records should support internal audit, regulatory examination, complaint handling, and post-incident review.

Pillar 6: Continuous monitoring and improvement. The organization should monitor data drift, model performance, alert volumes, case outcomes, segment disparities, user overrides, security incidents, and emerging typologies. Feedback from investigators should be incorporated through controlled updates. Periodic review should ask whether the tool remains necessary, whether simpler controls would now be sufficient, and whether benefits justify costs and risks.

For SMEs, implementation can be staged. Stage one involves digitizing customer files, standardizing risk factors, strengthening sanctions and adverse-media workflows, and measuring current alert outcomes. Stage two introduces limited AI assistance, such as document extraction, entity matching, or case prioritization, with human approval. Stage three adds behavioral or network analytics after sufficient data quality and governance capability exist. This staged approach reduces the risk of purchasing an advanced system before the organization can control it.



A minimum governance package for a smaller business should include an approved AI-use policy, an inventory of AI tools, a documented AML use case, vendor due diligence, a data and privacy assessment, pre-deployment testing, human-review requirements, incident reporting, performance monitoring, annual independent review, and a shutdown or fallback procedure. These controls are more important than whether the system uses a particular algorithm.

Implementation should also be supported by measurable performance indicators. Accuracy alone is insufficient because money laundering is rare, labels are imperfect, and the consequences of errors differ. Management should monitor alert volumes, alert-to-case conversion, investigator processing time, repeat alerts, data-quality exceptions, model overrides, and the proportion of cases escalated after human review. Where outcomes are available, the business may also examine whether the system improves the identification of previously unknown relationships or typologies. Metrics should be interpreted cautiously and should not reward staff merely for reducing alert numbers.

A practical pilot should compare the proposed AI-assisted process with the existing control. The pilot population should be sufficiently representative of customers and transactions, and it should include known difficult cases. Investigators should review outputs without being told that the new model is expected to perform better, reducing confirmation bias. The business should document what changed, what did not improve, and what new risks appeared. Deployment should occur only after accountable owners accept the residual risk and establish a rollback plan.

Incident management is another necessary component. An AI-AML incident may include unauthorized data disclosure, corrupted input data, unexplained changes in alert patterns, systematic misclassification, vendor outages, cyber manipulation, or the discovery that confidential information was entered into an unapproved generative AI service. The response plan should identify who can suspend the tool, how manual controls will continue, how affected decisions will be reviewed, and when regulators, customers, or other parties must be notified.

Training should be role-specific. Senior managers need to understand accountability, risk appetite, and resource implications. Compliance investigators need to understand the meaning and limits of model outputs. Technology personnel need to understand AML obligations, confidentiality, and evidentiary requirements. Internal auditors need sufficient access to assess governance and control effectiveness. Periodic exercises can test whether personnel recognize hallucinated content, synthetic identity indicators, misleading reason codes, and inappropriate automation.



7. Discussion

The analysis suggests that AI's greatest immediate value is not autonomous detection of money laundering but improved prioritization, consistency, and information organization. Money laundering is a legal and factual conclusion that often depends on context unavailable to a model. AI can identify patterns and inconsistencies, but investigators must interpret them in relation to the customer, product, jurisdiction, and known typologies. This supports an augmentation model in which technology expands human capacity.

The framework also shows that responsible AI and effective AML are mutually reinforcing. Data lineage supports both model validation and suspicious-activity documentation. Explainability supports both fairness and regulatory examination. Human oversight supports both legal accountability and error correction. Continuous monitoring supports both model-risk management and adaptation to criminal typologies. Accordingly, AI governance should not be treated as a separate technology project; it should be integrated into the AML control environment.

For managers, the principal implication is that technology procurement should follow risk definition rather than precede it. Organizations should resist the assumption that an AI label guarantees better compliance. A less complex, well-governed system may outperform a sophisticated model that lacks reliable data, user trust, and effective escalation. Management reporting should focus on outcomes and limitations rather than only the number of alerts processed.

For policymakers and supervisors, proportionate guidance is important. Smaller firms may avoid useful technology if expectations appear designed only for large banks. Conversely, vague encouragement of innovation may expose customers to untested systems. Supervisory guidance can help by clarifying expectations for validation, explainability, vendor management, privacy, and human oversight at different levels of risk and organizational complexity.

The framework has limitations. It is conceptual and does not quantify cost reductions, detection rates, or comparative performance. Legal obligations vary by jurisdiction and industry. The use of AI in a bank differs from its use by an accountant, real-estate company, online marketplace, or money-service business. Future research should test the framework through case studies, surveys of compliance professionals, controlled comparisons of rule-based and hybrid systems, and analysis of how model explanations affect investigator decisions.

8. Conclusion and Recommendations

Artificial intelligence can help businesses manage the scale and complexity of AML compliance by improving customer due diligence, monitoring, entity resolution, network analysis, and investigation support. Its value is strongest when it assists trained professionals

and operates within a clear risk-based framework. AI should not be used as a substitute for legal judgment, accountable governance, or adequate compliance staffing.

Businesses should adopt AI incrementally, beginning with defined problems and measurable controls. They should prioritize data quality, explainability, vendor transparency, human decision rights, and continuing validation. Confidential information should be protected, and generative-AI outputs should be verified before use. Senior management should receive reporting that includes limitations, overrides, drift, and adverse impacts, not merely efficiency statistics.

Regulators and industry bodies should provide proportionate guidance and encourage controlled testing, information sharing, and independent evaluation. Particular attention should be given to SMEs that face material laundering exposure but lack the resources of major financial institutions. A shared emphasis on governance can allow innovation without creating an unaccountable compliance system.

The central conclusion is that AI-enabled AML effectiveness depends less on the novelty of the algorithm than on the quality of the surrounding institutional controls. A responsible system is one in which risk is clearly scoped, data are governed, models are validated, humans retain authority, decisions are auditable, and performance is continuously reviewed. Under those conditions, AI can strengthen business integrity and contribute to a more resilient financial system.

Table

1

Responsible AI-AML Governance Framework

Pillar	Core requirement
1. Risk scoping	Define the AML problem, scope, decision, risk appetite, and proportionality.
2. Data governance	Control data quality, lineage, access, privacy, retention, and permitted use.
3. Model validation	Document logic, test performance, manage changes, and independently challenge results.
4. Human oversight	Assign decision rights, require review of consequential actions, and record overrides.
5. Auditability	Preserve evidence, reason codes, model versions, reviewer actions, and documentation.

6. Continuous monitoring	Monitor drift, outcomes, disparities, incidents, typologies, and system relevance.
--------------------------	--

Conflict of interest statement: The author declares no conflict of interest.

Funding statement: This research received no external funding.

AI-use disclosure: Generative AI was used as a drafting and language-support tool. The author reviewed, verified, and accepts responsibility for all arguments, citations, and final content.

References

- Bank for International Settlements. (2023). Project Aurora: The power of data, technology and collaboration to combat money laundering across institutions and borders. BIS Innovation Hub.
- Baxter, G., & Sommerville, I. (2011). Socio-technical systems: From design methods to systems engineering. *Interacting with Computers*, 23(1), 4-17.
- Desai, A., Kosse, A., & Sharples, J. (2024). Finding a needle in a haystack: A machine learning framework for anomaly detection in payment systems (BIS Working Papers No. 1188). Bank for International Settlements.
- European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*.
- Financial Action Task Force. (2020). Guidance on digital identity. Paris: FATF.
- Financial Action Task Force. (2021a). Opportunities and challenges of new technologies for AML/CFT. Paris: FATF.
- Financial Action Task Force. (2021b). Stocktake on data pooling, collaborative analytics and data protection. Paris: FATF.
- Financial Action Task Force. (2025). Horizon scan: Artificial intelligence and deepfakes—Impacts on money laundering, terrorist financing and proliferation financing. Paris: FATF.
- Financial Crimes Enforcement Network. (2018). Joint statement on innovative efforts to combat money laundering and terrorist financing. Washington, DC: U.S. Department of the Treasury.

- Financial Crimes Enforcement Network. (2024). FinCEN alert on fraud schemes involving deepfake media targeting financial institutions (FIN-2024-Alert004). Washington, DC: U.S. Department of the Treasury.
- Financial Stability Institute. (2019). Suptech applications for anti-money laundering (FSI Insights No. 18). Bank for International Settlements.
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). U.S. Department of Commerce.
- Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, Financial Crimes Enforcement Network, & National Credit Union Administration. (2018). Joint statement on innovative efforts to combat money laundering and terrorist financing.
- Organisation for Economic Co-operation and Development. (2019). Recommendation of the Council on artificial intelligence. Paris: OECD.
- Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3-38.
- United Nations Office on Drugs and Crime. (2011). Estimating illicit financial flows resulting from drug trafficking and other transnational organized crimes (Research Report). Vienna: UNODC.
- Wolfsberg Group. (2022). Principles for using artificial intelligence and machine learning in financial crime risk management.