



IOSCO Final Report - Supervisory Toolkit for AI Use in Capital Markets

October 2026

This Final Report provides IOSCO member authorities with practical, non-binding and non-prescriptive tools for **risk-based and proportionate supervision of AI systems used by supervised firms** in capital markets. It aims to help supervisors identify key considerations, adapt oversight to technological developments, and preserve market integrity, investor protection and financial stability.

The report builds on **IOSCO's previous work** from 2021¹ and 2025², however since the publication of these reports, **AI systems have continued to evolve significantly**, notably through the development and adoption of **Generative AI** (GenAI) and the emergence of **Agentic AI** techniques. These ongoing technological advances are broadening the range of AI applications in capital markets, offering potentially transformative benefits while also creating or exacerbating potential risks.

AI can offer significant **benefits** to capital markets, including greater efficiency, improved risk management, enhanced market analysis and forecasting, portfolio optimisation, fraud detection and prevention, compliance, and improved client service and engagement. At the same time, its increasing use introduces or amplifies **risks** relating to model and data quality, bias, transparency and explainability, governance and human oversight, third-party dependency and concentration, malicious use and operational vulnerability. The development of GenAI and Agentic AI may further increase these challenges, including through hallucinations, greater complexity and opacity, and, in the case of Agentic AI, increased autonomy and control risks.

As highlighted in the 2025 AI Report, these risks point to the need for **supervisory approaches that can adapt to technological developments**, particularly in light of the increasing complexity and opacity of certain AI systems, the continued evolution of AI use cases, and the potential for concentration and dependency on AI service providers.

The tools presented in this report are intended to address the **full lifecycle of individual AI systems** and to apply across **different types of AI systems**, ranging from those based on more traditional forms of Machine Learning (ML) to those incorporating GenAI and emerging Agentic AI techniques.

The toolkit is organised in **three complementary layers**:

1. Areas of supervisory consideration: The first layer sets out areas for supervisory consideration concerning the use of AI systems in capital markets, drawing on IOSCO's previous work on AI. It provides supervisors with a framework to identify areas that may require greater attention and to determine where supervisory resources can be deployed most effectively.

2. Tools for supervisory oversight of key areas: The second layer provides more detailed tools to assist supervisors in assessing four key areas: (i) Governance and Risk Management; (ii) Third-party and Outsourcing Risk Management; (iii) Disclosure; and (iv) Recordkeeping and Reporting.

3. Indicators and data sources: The third and final layer provides supervisors with suggested indicators to monitor the adoption and use of AI, together with a range of methods for gathering relevant information. These include on-site inspections, targeted surveys, regular supervisory dialogue, requests for documentation and data, and other mechanisms that may support supervisors in assessing firms' AI systems.

Through a survey, IOSCO has gathered **feedback from stakeholders on the supervisory toolkit** set out in this report. The responses will be considered as IOSCO continues to develop its work on AI. As a next step, IOSCO intends to review emerging industry practices in the supervisory areas identified in this report. It will continue to **coordinate efforts** relating to developments in the use of AI in capital markets and to **cooperate with other**

relevant international organisations, including the Financial Stability Board (FSB).

¹ <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf>

² <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf>

Pregunta

What AI use cases and developments does IOSCO identify in capital markets, and what benefits and risks accompany their adoption?

Respuesta

The report shows increasing AI use across capital markets. **Observed applications** include market analysis, forecasting and portfolio optimisation; support and risk-management functions; fraud detection and prevention; compliance; and client service and engagement. Previous IOSCO work also identified advisory and support services, client identification and monitoring, trading-algorithm selection, asset and portfolio management, robo-advising, algorithmic trading, investment research, sentiment analysis, surveillance and AML/CFT transaction monitoring. **GenAI** is increasingly moving from proofs of concept and pilots towards broader integration, while **Agentic AI** is still largely at the experimentation or proof-of-concept stage. Members also reported experimentation with AI used to assess other AI systems, synthetic data, and combinations of AI and blockchain. The **principal drivers** are operational efficiency and cost reduction, improved collection of high-quality data, client personalisation and experience, and risk management and fraud detection; AI is also reported to support compliance, reporting and detection of cyber threats. **Barriers** include governance gaps, reliance on third-party models/providers and shortages of skilled personnel. **Risks remain highly contextual** and include governance and oversight, model and algorithm development and testing, data quality and bias, transparency and explainability, outsourcing and concentration, malicious uses, ethical concerns and human-AI interactions. GenAI adds **uncertainty and hallucination risks**, while Agentic AI can **increase complexity** through planning, memory, external access and greater autonomy, with potential data, cybersecurity, operational, unexpected-behaviour and cascading-failure risks. IOSCO consequently supports **risk-based and proportionate supervision**, with greater supervisory attention for higher-risk or higher-impact uses.

Pregunta

What does the first layer of the toolkit - areas of supervisory consideration - cover?

Respuesta

The first layer gives supervisors a **framework for identifying where closer attention may be warranted and where supervisory resources can be deployed most effectively**. It builds on IOSCO's earlier AI work and focuses on the characteristics of the AI system, its use and deployment, its lifecycle, the degree of human oversight and its potential effects on investors, firms and markets. Supervisors are encouraged to assess the nature and complexity of the system, potential vulnerabilities in models or data, inappropriate use or cyber incidents, the effectiveness of risk mitigation and the firm's governance and risk-management arrangements. **Human oversight** can range from **human-in-control**, where the human may use or disregard the AI output; **human-in-the-loop**, where AI acts only after human approval; **human-on-the-loop**, where AI acts unless the human disapproves; and **human-out-of-the-loop**, where the system acts without human involvement. Impact considerations include investor protection, suitability, privacy and data protection, non-discrimination, the number and type of investors affected and the availability of redress, as well as potential disruption to firms' core activities and operational, legal or reputational consequences. Comprehensive **impact assessment** or **risk mapping** should support **proportionality**: higher-risk or higher-impact systems should receive more stringent examination, while lower-risk systems may receive more streamlined treatment. IOSCO also notes that systemically important institutions may warrant heightened consideration given their potential market-wide impact and that common dependence on third-party AI providers can create systemic vulnerabilities and cascading effects.

Pregunta

What tools does the second layer provide for supervisory oversight of governance and risk management, third-party and outsourcing risk, disclosure, and recordkeeping and reporting?

Respuesta

The second layer contains detailed supervisory tools and questions covering four areas. **Governance and risk management** should operate across the AI lifecycle, supported by documented policies, clear accountability, appropriate human oversight, escalation arrangements and audit trails. IOSCO highlights Board and senior-management oversight, firm-wide AI policies, AI inventories, risk and impact assessments, model validation and testing, ongoing monitoring for drift, degradation and bias, data governance, staff capability and training and practices specific to GenAI and other advanced systems. For **third parties and outsourcing**, firms should conduct risk-proportionate due diligence, assess transparency and fairness, validate third-party and open-source components, manage supply-chain and concentration risks, maintain contingency and exit arrangements and include appropriate contractual provisions on performance, data protection, audit and notification. The supervised firm retains responsibility, legal liability and accountability for outsourced tasks. **Disclosure** should give investors and clients meaningful information about material AI use, performance and limitations, relevant risks, governance and, where appropriate, client consent or opt-out arrangements and third-party dependencies. IOSCO also highlights misleading or promotional AI claims, or 'AI-washing'. **Recordkeeping** should preserve documentation and logs throughout the AI lifecycle, including system development and changes, model choices, AI-generated outcomes, human interventions, explanations and incidents. **Reporting** should provide supervisors with timely and sufficiently detailed information, including significant failures, biases, cybersecurity events, data breaches and other reportable AI incidents, while allowing existing reporting frameworks to be reviewed as AI use evolves.

Pregunta

What does the third layer - indicators and data sources for monitoring AI use - provide to supervisors?

Respuesta

The third layer supports ongoing supervisory monitoring by providing **indicators of AI adoption, use, dependencies and performance** and by identifying **ways to collect relevant information**. **Indicators** include the proportion of firms using AI; the proportion of use cases in development or production; AI research and development spending and patents; inventories of AI use cases and system types; information on underlying models and system attributes; dependence on and concentration among third-party providers; the frequency and severity of failures, errors, outages, cyber incidents, data breaches and misconduct; and model-performance measures such as prediction accuracy, error rates, drift and degradation. IOSCO also gives **sector-specific examples**, such as AI-driven mis-selling or unsuitable recommendations for broker-dealers, valuation and portfolio-drift measures for asset managers, and market-quality, surveillance, volatility and settlement indicators for exchanges and other market intermediaries. Supervisors may **obtain information** through supervisory reporting, examinations and inspections, structured engagement with firms, sector-wide surveys, publicly available information, real-time monitoring tools, AI-assisted analysis, information sharing among authorities, and sandboxes, innovation hubs and other engagement initiatives. The purpose is to identify adoption patterns and areas of heightened risk, prioritise supervisory engagement and detect emerging risks and vulnerabilities, while maintaining a risk-based and proportionate approach.

Links of interest:

[IOSCO Final Report - Supervisory Toolkit for AI Use in Capital Markets](#)

This document is an unofficial summary of a document published by IOSCO. It does not necessarily reflect CNMV's official position, its managers and directors' or the rest of its staff's on the matters addressed.

© CNMV. The contents of this publication may be reproduced, subject to attribution.