

AI GOVERNANCE IN EUROPEAN INVESTMENT MANAGEMENT

BARNA FLAVIA MIRELA

PROFESSOR, PHD., FACULTY OF ECONOMICS AND BUSINESS ADMINISTRATION, WEST
UNIVERSITY OF TIMISOARA, ROMANIA

e-mail: flavia.barna@e-uvt.ro

ȚĂRANU ANCA

MASTER'S STUDENT, FACULTY OF ECONOMICS AND BUSINESS ADMINISTRATION, WEST
UNIVERSITY OF TIMISOARA, ROMANIA

e-mail: anca.taranu91@e-uvt.ro

Abstract

With the increasing adoption of artificial intelligence by investment managers in Europe, some questions arise that go beyond predictive performance. Investment firms will need to consider not only whether an AI system is caught by the scope of the EU AI Act, but also how its use interacts with existing governance, conduct, algorithmic trading and operational resilience requirements. The paper discusses regulatory constraints that apply to AI systems used for investment research, portfolio construction, risk management and algorithmic trading. The study is founded on a systematic documentary analysis of the EU AI Act, MiFID II, RTS 6, DORA and recent ESMA publications in order to distinguish between formal legal obligations and more general supervisory expectations and model governance. This demonstrates that portfolio selection models are not themselves high-risk systems under the AI Act and that explainability is not a uniform legal requirement for all investment models. But transparency and auditability still matter for internal validation, human oversight, model-risk management and supervisory accountability. 'This paper introduces a proportionate governance framework with three dimensions: the role of the AI system in the investment process, the degree of autonomy of the AI system and the potential impact on clients, firms and market integrity'. The framework connects these dimensions to differentiated requirements for documentation, validation, explainability, monitoring, human intervention, and third party oversight. The analysis will provide an operational framework for European asset managers to balance technological innovation with accountability and regulatory compliance.

Keywords: artificial intelligence, investment management, AI governance, EU AI Act, model explainability, algorithmic trading

JEL Classification: G11 ; G18 ; G28 ; O33

1. Introduction

Investment firms and asset managers are increasingly embedding artificial intelligence into an ever-growing range of activities. Current uses include document processing, research assistance, data extraction, investment signal generation, portfolio optimisation, risk monitoring, client engagement, and algorithmic execution. The purposes, the level of autonomy and the potential impact of these applications are very diverse. Governance issues arise not from an internal tool that summarises research but from a system that recommends portfolio changes or automatically generates and submits trading orders.

The integration of artificial intelligence into investment processes has significantly accelerated in the last decade. In a Mercer survey of 131 asset managers around the globe in February 2026, 55% of respondents said AI had been embedded into at least one investment process and 91% said they planned to grow their use of AI in the next 12 months. But only 8 per cent reported any measurable improvement in investment returns. The benefits cited by respondents were largely operational in nature, such as greater efficiency and faster or better insights. The disparity between adoption and measurable investment performance suggests that the practical

value of AI cannot be assessed solely by implementation rates or technological capabilities (Mercer, 2026).

The adoption–performance gap in Europe has a particular regulatory and supervisory dimension, something that is less pronounced in much of the international literature on AI-enabled investment management. While the EU AI Act provides a comprehensive horizontal framework for the development and use of AI systems, sector-specific requirements still apply through MiFID II, RTS 6 and DORA. European evidence also suggests a differentiated pattern of uptake. The ESMA survey of securities-market entities shows that the applications of AI are still concentrated in the area of internal support functions, and less so in portfolio optimisation and algorithmic trading (ESMA, 2026a). Separately, Bagattini and Piazza (2025), analysing approximately 800,000 regulatory and marketing documents issued by 44,000 EU funds, identified only 145 funds that disclosed the use of AI or machine learning in their investment strategies. These funds represented approximately 0.1% of UCITS assets under management and did not exhibit significantly different net performance from comparable actively managed funds.

Industry and academic discussions often oversimplify the regulatory treatment of these systems. In particular, model interpretability is in some cases put forward as a uniform duty set out in Regulation (EU) 2024/1689, also called the EU AI Act. That's not the whole interpretation. The AI Act's risk-based structure and requirements depend on the intended purpose and legal classification of a system. Not all portfolio selection and investment research is automatically a high risk use case. At the same time, the fact that it is not classified as high risk does not mean that effective governance is not needed. Depending on DORA, MiFID II, RTS 6 and ESMA supervisory expectations, investment firms may need to understand, test, monitor, document and control the systems upon which they rely.

The governance problem thus created is broader than a formal classification exercise. Investment firms need to identify which regulatory instruments apply, distinguish legal obligations from supervisory expectations and good practice and calibrate their controls to the actual role of AI in the investment process. A model that is technically complex might only serve as an internal research tool, while a comparatively simple model might have a material impact on client portfolios or on market execution. Therefore, technical complexity alone cannot be used to determine the appropriate level of oversight.

The paper poses four research questions. First, how does the EU AI Act categorise AI systems used in investment management, portfolio construction and algorithmic trading? Second, what transparency, governance and explainability expectations stem from the interplay of the AI Act, MiFID II, RTS 6, DORA and ESMA guidance? Third, how should governance requirements be calibrated to the role, the autonomy and the potential impact of an AI system? Fourth, what practical framework can European asset managers use to balance innovation, accountability, auditability and supervisory compliance?

The paper does not assess whether machine-learning models produce better investment returns. Its contribution is quite normative and organisational. It produces a proportional governance framework linking the function, autonomy, and potential consequences of an AI system with certain requirements for documentation, validation, explainability, human oversight, monitoring, and third-party control. The remaining part of the paper is organised as follows. We introduce the research design, classify AI use cases in investment management, analyse the relevant European framework, discuss the role and boundaries of explainability, and develop the proposed governance matrix.

2. Research design

The research is grounded in a systematic review of regulations and documents. This method is appropriate because the research aims to explain the interaction of several legal and supervisory instruments and not to estimate the financial performance of a certain algorithm. The documentary corpus includes binding European legislation, delegated regulations, official supervisory publications and selected frameworks of international governance and academic contributions on interpretable machine learning and model risk.

The core legal corpus includes the EU AI Act, Directive 2014/65/EU on markets in financial instruments (MiFID II), Commission Delegated Regulation (EU) 2017/589 (RTS 6) and Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). The supervisory corpus comprises ESMA’s public statement on the use of AI in retail investment services in 2024, its analysis of the use of AI in EU investment funds in 2025, its analysis of AI adoption in securities markets in 2026, and its supervisory briefing on algorithmic trading in 2026. International reference points include the IOSCO report on artificial intelligence and machine learning used by market intermediaries and asset managers, the Financial Stability Board’s analysis of artificial intelligence in financial services and the NIST AI Risk Management Framework. Contextual industry evidence is provided by Mercer’s 2026 global survey of 131 asset managers, which is used to compare organisational adoption with disclosed fund-level use and measurable investment outcomes.

Each document was analysed along nine analytical dimensions: scope and intended use; legal risk classification; transparency and explainability; human oversight; model testing and validation; data governance; documentation and traceability; third-party dependency; and continuous monitoring. The identified requirements and expectations were matched to five categories of AI use in investment management. The final synthesis is based on a principle of proportionality: the intensity of governance should increase with the system’s decision-making role, degree of autonomy, reversibility of outcomes and potential impact on clients, the firm and market integrity.

Table 1. Documents included in the core regulatory and supervisory corpus

Document	Primary focus	Relevance to AI governance
EU AI Act (2024)	Risk-based AI regulation	Classification, transparency, oversight, documentation
MiFID II and RTS 6	Investment services and algorithmic trading	Firm accountability, testing, controls, monitoring
DORA (2022)	Digital operational resilience	ICT risk, resilience, incidents, third parties
ESMA (2024)	AI in retail investment services	Client interest, governance, data quality, oversight
ESMA (2025)	AI in EU investment funds	Adoption, use cases, fund-level evidence
ESMA (2026a)	AI adoption in securities markets	Use cases, risks, infrastructure, providers
ESMA (2026b)	Algorithmic trading supervision	AI in trading, validation, explainability, pre-trade controls

Source: authors’ synthesis.

3. AI Use Cases for Investment Management

AI should not be seen as single use case in investment management. The same technical architecture can be employed for activities with very different legal and operational effects. A useful classification starts with the function performed by the system and the degree of influence or determination of its outputs on a financial decision.

3.1. Support and Operations Applications

The lowest level of direct investment impact is the use of AI for drafting, summarising, translation, document classification, coding assistance, data extraction and internal search. These internal support applications are the most reported AI use cases according to ESMA's 2026 survey (ESMA, 2026a). Their key risks are confidentiality, personal data, hallucinated or incomplete outputs, unauthorised disclosure, cyber security and over-reliance on external providers. Human review is typically available before an output affects a client or portfolio.

3.2. Research & Signal Generation for Investment

A second category consists of systems that process financial statements, news, market data and alternative data or analyst reports, and generate scores, forecasts or investment signals. These systems may not be the ultimate decision makers, but their output can affect the information set used by portfolio managers. Therefore, governance should include data provenance, feature construction, temporal validity, bias, model drift, output stability and the ability of investment professionals to question a recommendation. Formal human approval is not enough if users do not have the information or competence to challenge the model.

3.3. Portfolio Formation and Risk Management

AI systems can also suggest asset allocation, security selection, position limits, hedging decisions or rebalancing. Their output can directly impact portfolio composition, client risk, liquidity, concentration, and suitability. And, although the portfolio manager has the final say, always heeding the automated recommendations may lead to automation bias and turn nominal oversight into a routine approval. Controls should include independent validation, benchmarking, sensitivity analysis, override procedures, performance monitoring and defined escalation thresholds.

3.4. Trading and Execution Algorithms

The most immediate market impact comes from embedding AI into an algorithmic trading system that decides elements of orders with limited or no human intervention. Here, the applicable framework is broader than the AI Act. Under MiFID II, Article 17 and RTS 6 set out requirements for algorithmic trading in terms of organisation, testing, control, resilience and documentation. Companies require effective control over their algorithms, including the outsourcing of development or infrastructure. Pre-trade controls, real-time monitoring, business continuity arrangements and the ability to suspend a system are all important, because errors can propagate at market speed.

4. THE EUROPEAN REGULATORY SYSTEM

4.1. EU AI Act: Scope and Risk Classification

The AI Act establishes harmonised rules for AI systems that are placed on the market, put into service or used in the European Union. It is risk based in structure. Article 5 prohibits certain practices that are deemed unacceptable. High-risk systems must comply with a host of obligations on risk management, data governance, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness and cybersecurity. Article 50 sets out some specific

transparency obligations for some systems, including some systems that directly interact with natural persons or generate or manipulate content.

In the financial services sector, a single high risk domain is not considered. Specific financial use cases are listed in Annex III, especially AI systems used for the assessment of creditworthiness of natural persons or the establishment of their credit score, and certain systems used for risk assessment and pricing in life and health insurance. Investment research, portfolio selection and asset allocation are not automatically included in the categories below. Thus, it is not the case that all AI models which inform an investment decision need to comply with the high-risk transparency and human-oversight requirements established by the AI Act.

This does not mean that the Act is irrelevant for investment firms. A firm can be a provider or a deployer of an AI system and needs to identify what provisions apply according to the intended purpose and deployment context of the system. Duties of AI-literacy and other general provisions may also be relevant. Also, a system developed for internal support may be modified or incorporated into a more impactful process. Classification therefore needs to be documented and reviewed whenever the specific purpose, the users and their autonomy, and the context of deployment change. Figure 1 summarises the difference between the AI Act risk framework and the governance implications of model interpretability in investment applications.




Figure 1. EU AI Act risk framework and governance implications for investment models. Source: authors, based on Regulation (EU) 2024/1689 and ESMA (2026b).

4.2. MiFID II, RTS 6 and Investment Firm Accountability

MiFID II offers a sector-specific tier of accountability. In its 2024 public statement, ESMA stressed that investment firms continue to be responsible for compliance when they use AI for the provision of investment services. Relevant duties include acting in good faith, fairly and professionally in the best interests of clients, maintaining appropriate organisational arrangements, assuring the quality of information and managing conflicts of interest. The use of external or generative AI tools does not shift these responsibilities to a vendor (ESMA, 2024).

Article 17 of MiFID II and RTS 6 set out more detailed requirements for algorithmic trading. Governance arrangements should define accountability, separate relevant tasks, ensure compliance personnel understand how trading algorithms work, and include provisions for testing, validation, monitoring, business continuity and effective pre-trade controls. The 2026 ESMA supervisory briefing explicitly states that AI-based algorithmic trading is not, as of now, identified as a high-risk use case under the AI Act. But firms must assess the relevant provisions of the AI Act and be able to explain how AI affects the decision-making of their algorithms in the MiFID II

and RTS 6 governance framework (ESMA, 2026b). First, you must distinguish between portfolio management and algorithmic trading. An algorithmic trading system is not by definition a model that gives a monthly portfolio suggestion. Classification changes when the system determines parameters of orders automatically or initiates execution with minimal human intervention. As a result, governance assessments should follow the complete decision chain from data ingestion and signal generation to portfolio approval, order generation, routing, execution and post-trade monitoring.

4.3. DORA and Third Party Operational Risk

DORA contributes the operational-resilience perspective. AI systems deployed by financial entities depend on data pipelines, computing infrastructure, cloud services, application programming interfaces, model providers and other ICT third parties. Any one of these components failing can degrade investment operations even though the model itself is working as designed. DORA therefore increases the need for ICT risk management, incident handling, resilience testing, business continuity, contractual controls and oversight of third-party arrangements.

Third party concentration is particularly relevant to AI. ESMA's 2026 evidence points to a strong reliance on a handful of commercial infrastructure and model providers (ESMA, 2026a). Companies need to know where models and data are hosted, which subcontractors are used, how model updates are introduced, whether outputs are reproducible, and how service termination or provider failure will be managed. There must be contractual access to documentation, logs, testing evidence and information about incidents. But no contractual language can substitute for internal competence and good control.

4.4. Supervisory Evidence on AI Adoption

The regulatory analysis should also be understood in the context of the real pattern of adoption of AI. ESMA's survey of the securities market and its separate analysis of EU investment-fund disclosures are different populations and should not be aggregated as a single adoption indicator. In the survey on the securities market, about 45% of responding entities indicated that they did not have an AI use case in production, development or experimentation, while reported applications were mainly focused on internal support tasks (ESMA, 2026a).

The fund-disclosure analysis identified 145 funds claiming to use AI or ML, representing approximately 0.1% of UCITS assets under management, with an estimated 30% of these funds being classified as truly AI-driven by number. Their net performance did not differ significantly from that of similar actively managed funds (Bagattini and Piazza, 2025). Figure 2 shows the two evidence bases separately and does not equate firm-level adoption with fund-level disclosure.

The Mercer and ESMA findings should be interpreted as complementary rather than directly comparable. The Mercer survey captures self-reported organisational adoption among global asset managers, whereas the ESMA fund analysis identifies explicit references to AI or machine learning in regulatory and marketing documents. Together, they indicate that organisational experimentation and operational integration are advancing more rapidly than the use of AI as a disclosed and performance-differentiating investment strategy.




Figure 2. AI adoption in EU securities markets and EU investment funds: distinct evidence bases. Source: ESMA (2026a) for the securities-market entity survey; Bagattini and Piazza (2025) for the fund-disclosure analysis.

5. Explainability as a governance tool

5.1. Direct & Post-Hoc Interpretability

Interpretability is the degree to which one can understand the rationale for model inputs and the resulting outputs. Some models provide relatively simple explanations. Coefficients in an appropriately specified regression model indicate the direction and magnitude of associations. It can be traced from output to input conditions by a shallow decision tree. More complex ensembles and neural networks often need additional tools to summarise global behaviour or explain individual predictions.

Post-hoc methods, such as SHAP, attribute a prediction to input features and can be useful for diagnosis, communication and challenge (Lundberg and Lee, 2017). However, they do not make a complex model intrinsically interpretable. Another model or approximation of behaviour is an explanation. It may be sensitive to the reference distribution, feature dependence, local neighbourhood, or method chosen. Rudin (2019) therefore argues that in high stakes settings, where interpretable models can achieve adequate performance, we should consider using those models before we rely only on explanations of black-box systems.

5.2. Purpose and Audience of Explainability

No one explanation is suitable for all purposes. The economic drivers of a recommendation may need to be understood by a portfolio manager. Model validators need architectural details, assumptions, training data, stability and sensitivity information. Compliance people need to know if the process complies with legal and organisational requirements. Senior management needs a single view of limits and material risks. Supervisors may require documentation, test reports, decision logs and a clear story of how the system is controlled. Clients should be given clear explanations, which are easy to understand, without going into unnecessary technical detail.

The design of explainability is thus audience and purpose dependent. A technically sophisticated feature attribution chart may be good enough for a data scientist, but not for board oversight. A simplified story might be sufficient for communicating with the client but not for validation. Good governance is a combination of different types of evidence: model documentation, data lineage, performance and stability tests, local and global explanations, override records and incident log, with clear accountability.

5.3. The Human Oversight Dilemma

Human oversight can only work if the human in charge has enough time, competence, authority, and information to intervene. Previous studies on automation have demonstrated that nominal human involvement may be accompanied by automation bias, loss of situational awareness, and difficulty taking over control after prolonged reliance on automated systems (Bainbridge, 1983; Parasuraman, Sheridan and Wickens, 2000). If a recommendation is accepted automatically or the basis for challenging it is not available, requiring a portfolio manager to approve every recommendation does not constitute meaningful oversight. Governance should define the circumstances in which human review is required, the information provided to the reviewer, the timing of intervention, the ability to override or suspend the system, and what happens if there are repeated overrides. Exceptionally low override rates might suggest over-reliance, whereas persistently high rates might imply that the model is not appropriate for its intended use.

6. A Framework of Ratio Governance

The proposed framework considers governance intensity as a function of three dimensions. The first is the role of the system: internal support, analysis, recommendation, portfolio decision or execution. The second is autonomy, or whether a human reviews each output, monitors the system at an aggregate level, or is not in the immediate decision loop at all. The third is the potential impact – the scale, speed, reversibility and consequences of an error for clients, the firm and market integrity. These dimensions are to be measured together.

Note that levels of governance in Table 2 are not legal categories under the AI Act. A tool of internal control that helps firms proportion resources for validation, documentation and oversight. The tiering may vary depending on the system used. For example, a research model used for generating ideas may be subject to tighter controls if the output is subsequently directly linked to portfolio rebalancing.

Table 2. Proportionate AI-governance matrix for investment management

Governance tier	Typical use	Impact and autonomy	Minimum governance controls
1. Support	Drafting, summarization, translation, coding, document processing	Indirect investment impact; human review before use	Access and confidentiality controls; output verification; approved tools; AI literacy; vendor assessment
2. Analysis	Research, data extraction, forecasting, scoring, signal generation	Influences the information set used by professionals	Data lineage; benchmark testing; documented limitations; drift monitoring; competent human challenge
3. Recommendation	Portfolio proposals, asset allocation, risk limits, rebalancing recommendations	Material influence on portfolios or clients; human approval retained	Independent validation; explainability; sensitivity analysis; override procedures; approval and escalation thresholds
4. Decision	Automated portfolio construction or rebalancing within delegated limits	High autonomy and direct financial impact	Formal model inventory; independent approval; continuous monitoring; hard limits; fallback process; audit trail

Governance tier	Typical use	Impact and autonomy	Minimum governance controls
5. Execution	Algorithmic order generation, routing, execution, or high-speed adjustment	Immediate market impact; limited time for intervention	MiFID II/RTS 6 controls; pre-trade limits; real-time monitoring; kill capability; resilience testing; incident management

Source: authors' framework. The tiers indicate internal governance intensity and do not reproduce the legal risk classes of the EU AI Act.

6.1. Core Principles of Governance

The framework is built on eight principles. First, proportionality means controls should be proportionate to use, autonomy and impact. Second, accountability means a named business owner and this accountability remains with the investment firm. Third, human oversight has to be meaningful and it requires competence, authority and the ability to intervene in practice. Fourth, traceability includes tracking data, model versions, decisions and overrides. Fifth, independent validation separates the approval of the model from the development of the model. Sixth, performance degradation, drift, and changing market dynamics are monitored in production. 7. Explainability by audience: Personalise evidence for portfolio managers, validators, management, clients and supervisors. 8. Third-party governance prevents loss of internal control due to outsourcing.

These principles are consistent with broader frameworks of model risk. The NIST AI Risk Management Framework organises AI governance along the functions of govern, map, measure, and manage (NIST, 2023). IOSCO highlights the importance of governance, testing, monitoring, data quality, disclosure, and outsourcing controls for market intermediaries and asset managers (IOSCO, 2021). The proposed matrix translates these general principles into the functional sequence of investment management.

6.2. Governance Life Cycle

Proportional governance should cover the whole system life cycle. The firm should determine the intended purpose, prohibited uses, decision owner, affected stakeholders, data requirements and the applicable regulatory perimeter before developing or acquiring the firm. Development was supposed to include controlled data lineage, reproducible training, benchmark models, robustness tests, and documented limitations. Independent validation should include conceptual soundness, implementation, performance, stability, explainability and operational integration.

Deployment should be subject to approved thresholds, user training, access controls, monitoring metrics, escalation rules, and fallback arrangements. You need to watch data quality, output distributions, drift, exceptions, overrides, incidents and third party changes. Re-assessment and re-validation should be conducted as appropriate to changes in the materiality of the data, model architecture, intended purpose or execution logic. Retirement should involve archiving of records, and a managed hand-over to replacement processes.

7. Conclusions

Our analysis yields three main implications. First, regulators cannot simply consider that the AI used in investment management is in the high risk categories of the AI Act. Material conduct, model, operational, liquidity, concentration or market integrity risks may also be generated by a portfolio-selection system outside the scope of Annex III. So, sectoral rules and supervisory expectations are still at the center. Second, explainability should not be treated as a binary property or a universal compliance certificate. Interpretable models have important advantages, but

transparency alone does not guarantee good data, stable performance, proper incentives or good governance. That is not to say that complex models are automatically out of bounds or inappropriate. The question is whether the firm can justify the complexity of the model, validate the model's behaviour, understand the model's limitations, monitor the model's operation and intervene effectively. Third, governance must be throughout the whole decision chain. Research, portfolio construction and execution are often treated as separate systems, which can mask cumulative risk. A low impact signal model can be very important when automatically combined with portfolio rules and order execution logic. This transition is captured in the proposed role-autonomy-impact assessment and guarantees that governance is not only defined by the name or architecture of an algorithm. The framework also has implications for organising. AI governance cannot just be the job of data scientists or compliance departments.

Investment professionals, risk management, compliance, information security, legal functions, internal audit, procurement and senior management need to work together for effective oversight. AI literacy needs to be tailored to responsibilities: portfolio managers need to challenge outputs, validators should have access to technical evidence and management needs to understand aggregate exposure and residual risk.

European investment management needs layered, proportionate AI governance. Portfolio-selection and investment-research models are not, per se, high-risk systems under the EU AI Act. The EU AI Act does not impose a uniform explainability obligation for all investment models. But they may also create significant model, conduct, operational and market risks. Therefore, MiFID II, RTS 6, DORA and ESMA supervisory expectations still provide relevant guidance on the way AI should be governed in investment firms.

The level of required governance should be dependent on the role of the system, the extent of its autonomy and its potential impact. Internal support tools require data controls, access controls, accuracy controls and third party use controls. Investment advice-generating systems need to be better validated, documented, challenged and monitored. Portfolio decision making or order execution systems should have independent validation, real-time controls, clear authority to intervene, arrangements for resilience and complete audit trails.

The best way to think about explainability is as part of governance, not as a broad legal standard. It should be tailored to the target audience and decision context and presented with evidence on data quality, performance, robustness, stability and control of operations. Competent and authoritative substantive human oversight should not be boiled down to mere formal approval.

The proposed framework provides an operational framework for European asset managers to translate the fragmented regulatory demands into differentiated controls. The simple idea is that governance should be based on the use, autonomy and impact, not the complexity of technology. The framework enhances innovation, accountability, auditability, client protection and market integrity.

References

- [1] Bagattini, G., and Piazza, F. (2025). *Artificial intelligence in EU investment funds: Adoption, strategies and portfolio exposures*. ESMA TRV Risk Analysis, ESMA50-43599798-9923, European Securities and Markets Authority, Paris.
- [2] Bainbridge, L. (1983). *Ironies of automation*. *Automatica*, 19(6), 775–779. [https://doi.org/10.1016/0005-1098\(83\)90046-8](https://doi.org/10.1016/0005-1098(83)90046-8).
- [3] European Banking Authority. (2020). *Report on Big Data and Advanced Analytics*, EBA/REP/2020/01, Paris.
- [4] European Commission. (2016). *Commission Delegated Regulation (EU) 2017/589 supplementing Directive 2014/65/EU with regard to regulatory technical standards specifying the organisational requirements of investment firms engaged in algorithmic trading*. Official Journal of the European Union, L 87, 417–448.

- [5] European Parliament and Council. (2014). *Directive 2014/65/EU on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU*. Official Journal of the European Union, L 173, 349–496.
- [6] European Parliament and Council. (2022). *Regulation (EU) 2022/2554 on digital operational resilience for the financial sector*. Official Journal of the European Union, L 333, 1–79.
- [7] European Parliament and Council. (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. Official Journal of the European Union, L 2024/1689.
- [8] European Securities and Markets Authority. (2024). *Public Statement on the Use of Artificial Intelligence in the Provision of Retail Investment Services*. ESMA35-335435667-5924, Paris.
- [9] European Securities and Markets Authority. (2026a). *AI adoption and trends in securities markets: EU evidence*. ESMA TRV Risk Analysis, ESMA50-481369926-30599, Paris.
- [10] European Securities and Markets Authority. (2026b). *Supervisory Briefing on Algorithmic Trading in the EU*. ESMA74-1505669079-10311, Paris.
- [11] Financial Stability Board. (2017). *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications*. Basel.
- [12] International Organization of Securities Commissions. (2021). *The Use of Artificial Intelligence and Machine Learning by Market Intermediaries and Asset Managers*. Final Report, FR06/21, Madrid.
- [13] Lundberg, S.M., and Lee, S.I. (2017). *A unified approach to interpreting model predictions*. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- [14] Mercer. (2026). *How Artificial Intelligence Is Shaping Asset Management*. Mercer Artificial Intelligence Asset Manager Survey, February 2026. <https://www.mercer.com/en-us/about/newsroom/how-artificial-intelligence-is-shaping-asset-management/>
- [15] National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1, Gaithersburg, MD. <https://doi.org/10.6028/NIST.AI.100-1>.
- [16] Parasuraman, R., Sheridan, T.B., and Wickens, C.D. (2000). *A model for types and levels of human interaction with automation*. *IEEE Transactions on Systems, Man, and Cybernetics – Part A*, 30(3), 286–297. <https://doi.org/10.1109/3468.844354>.
- [17] Rudin, C. (2019). *Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead*. *Nature Machine Intelligence*, 1, 206–215. <https://doi.org/10.1038/s42256-019-0048-x>.