ISSN: 3033-4136
DOI: https://doi.org/10.36399/gla.pubs.382504

AI Governance *after* MiFID II: Beyond (Mere) Technological Neutrality?

FRIL
Financial Regulation Innovation Lab

FinTech Scotland®

University of Glasgow

University of Strathclyde Glasgow

AI Governance *after* MiFID II: Beyond (Mere) Technological Neutrality?

Alessio Azzutti

School of Law, University of Glasgow, and Financial Regulation Innovation Lab, Stair Building, 5 – 9 The Square, Glasgow G12 8QQ UK

19 March 2026

We acknowledge funding from Innovate UK, award numbers 10055559 and 10157846.

Corresponding author:

Email: Alessio.Azzutti@glasgow.ac.uk

Financial Regulation Innovation Lab

Who are we?

The Financial Regulation Innovation Lab (FRIL) is an industry-led collaborative research and innovation programme focused on leveraging new technologies to respond to, shape, and help evolve the future regulatory landscape in the UK and globally, helping to create new employment and business opportunities, and enabling the future talent.

FRIL provides an environment for participants to engage and collaborate on the dynamic demands of financial regulation, explore, test and experiment with new technologies, build confidence in solutions and demonstrate their ability to meet regulatory standards worldwide.

FRIL is part of the Glasgow City Region Innovation Accelerator programme, funded through Innovate UK on behalf of UK Research and Innovation. The Innovation Accelerator programme is investing £130 million in 26 transformative R&D projects to accelerate the growth of three high-potential innovation clusters, including the Glasgow City Region.

What is Actionable Research?

FRIL will integrate academic research with an industry relevant agenda, focused on enabling knowledge on cutting-edge topics such as generative and explainable AI, advanced analytics, advanced computing, and earth-intelligent data as applied to financial regulation. The approach fosters cross sector learning to produce a series of papers, actionable recommendations and strategic plans that can be tested in the innovation environment, in collaboration across industry and regulators.

Locally-led Innovation Accelerators delivered in partnership with DSIT, Innovate UK and City Regions

Department for Science, Innovation & Technology

UKRI Innovate UK

GLASGOW CITY REGION

FRIL White Paper Series

AI Governance *after* MiFID II: Beyond (Mere) Technological Neutrality?

Alessio Azzutti

University of Glasgow

19 March 2026

Abstract

This article examines the evolving intersections between artificial intelligence (AI) and EU financial regulation, focusing on the Markets in Financial Instruments Directive II (MiFID II). Grounded in the principle of technological neutrality, MiFID II seeks to enhance investor protection, safeguard market integrity, and ensure that innovation develops within competitive and well-regulated markets across the Union. The article argues, however, that while this neutrality renders the framework functionally enabling, it also leaves it normatively silent in the face of the distinctive and evolving risks introduced by financial AI. As AI applications become increasingly heterogeneous—both across the financial functions in which they are deployed and in their underlying lifecycles and value chains—MiFID II's activity-based logic increasingly struggles to accommodate their diverse and evolving risk profiles. Reflecting the EU's broader shift toward risk-based AI governance, the article outlines an initial taxonomy of financial AI applications designed to guide the proportionate alignment of regulatory obligations with AI-related risks, thereby supporting the continued adaptability, coherence, and future-proofing of EU financial services law.

Keywords: Artificial Intelligence; MiFID II; Technological Neutrality; Risk-based Regulation; AI Governance.

Table of Contents

1. Introduction

Artificial intelligence[1] (AI) has moved from the periphery to the core of European Union (EU) financial sector. Once limited to rule-based data analytics and routine process automation, AI—particularly its subfields of Machine Learning (ML) and, more recently, Generative AI (GenAI)—now permeates nearly every stage of the financial services lifecycle. According to the European Banking Authority (EBA), over 90 per cent of EU banks already deploy AI in some form.[2] AI systems now underpin a broad range of functions, including customer profiling, risk modelling, credit scoring, anti-money laundering (AML) surveillance, financial crime detection, customer support, regulatory reporting, and business process optimisation.[3] The European Central Bank (ECB) likewise reports broad use of 'conventional' AI tools and growing experimentation with GenAI across client-facing and internal operations.[4] Across capital markets, the European Securities and Markets Authority (ESMA) observes growing AI adoption, with firms applying it mainly to research, risk management, compliance, and post-trade analytics, while fully AI-based services remain limited.[5] In EU investment funds, explicit references to AI have risen too, with about one-third identifying it as a core investment driver.[6] Despite these trends, supervisors only have limited visibility over actual AI deployment by firms, as demonstrated by a recent survey launched by ESMA to gather more systematic information on AI strategies, investment levels, operational use cases, and governance.[7] These findings reflect a wider acceleration of investment in financial AI. Although EU private investment in FinTech and AI remains modest compared with the United States and the United Kingdom, recent data indicate a sharp rise—partly driven by the spread of AI sandboxes across Member States—as financial institutions and FinTechs pursue competitiveness and efficiency gains.[8] Yet the growing reliance on AI also introduces new prudential, conduct, and financial stability risks. European Supervisory Authorities (ESAs) have warned of vulnerabilities related to model opacity and explainability, data quality and integrity, cybersecurity, third-party dependency and concentration, and emergent behaviour.[9] These concerns echo those of global standard-setters, who warn that AI may open new channels of systemic risk.[10]

Despite these developments, the cornerstone of EU financial markets law—the Markets in Financial Instruments Directive II (MiFID II)[11]—was crafted in a different socio-technological era. Grounded in the principle of technological neutrality, MiFID II applies its regulatory objectives across financial services regardless of delivery method, with specific provisions for algorithmic and high-frequency trading (HFT). This functional approach ensures that investor protection, conduct of business, and organisational requirements apply regardless of whether financial services are provided by humans, assisted by AI, or are fully AI-driven. Yet this regulatory model increasingly shows its limits, as AI-related risks are heterogeneous and

[1] For the purposes of this article, the term "AI" refers to computational systems—typically software-based—designed to perform cognitive or analytical tasks that would ordinarily require human intelligence, such as learning from data, recognising patterns, or making decisions under uncertainty.
[2] EBA [18].
[3] EBA [21].
[4] *Leitner, G. et al.* [38].
[5] ESMA [32].
[6] ESMA [28].
[7] See *Naeem/Treacy* [41].
[8] AFME [1], p. 25.
[9] E.g., ESAs [26].
[10] E.g., IOSCO [36]; FSB [34]
[11] Directive 2014/65/EU of the European Parliament and of the Council of 15.5.2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (recast) [2014] OJ L 173/349.

context-specific. While technological neutrality has long been praised for balancing regulatory objectives and fostering innovation, its capacity to "future-proof" financial legislation is now open to question.[12] The pace and nature of AI innovation challenge the assumption that existing, technology-neutral obligations are sufficient to address emerging, technology-specific risks. MiFID II's design thus yields broad coverage but little differentiation. Meanwhile, the EU has begun to articulate a cross-sectoral response to these challenges through the Artificial Intelligence Act (AI Act).[13] Adopted in March 2024 and entering into force in 2025, the AI Act established the first horizontal framework for AI governance, classifying systems by risk level and imposing proportionate obligations throughout the AI lifecycle and value chain. Although conceived as a horizontal instrument, it explicitly encompasses certain financial AI applications—such as credit scoring and creditworthiness assessment—within its high-risk category.

Amid the ongoing evolution of EU financial and AI regulation, this article advances two interrelated claims. First, MiFID II is *functionally enabling yet normatively silent* toward AI. Its neutrality promotes flexibility and facilitates technological innovation but downplays the need for differentiated regulatory treatment across different applications. Second, given the heterogeneity of AI technologies and their associated risks to consumers, investors, firms, and markets, EU financial law should evolve toward a more pronounced risk-based model of AI governance. Rather than abandoning neutrality altogether, MiFID II should evolve toward a more risk-proportionate approach, maintaining flexibility while tailoring obligations to the actual risks posed by different AI uses. These claims structure the analysis that follows, which proceeds in three steps. Section 2 maps the current treatment of AI under MiFID II across the financial services lifecycle, distinguishing four key domains: (i) consumer-facing, (ii) market-facing, (iii) firm-internal, and (iv) supervisory applications. Section 3 examines the limits of technological neutrality in governing AI within investment services and proposes a risk-based framework to introduce greater proportionality into MiFID II's architecture. Section 4 concludes.

2. The Legal Treatment of Artificial Intelligence under MiFID II

AI systems increasingly support, or even perform, core functions regulated under MiFID II. Designed in the aftermath of the 2008 financial crisis, the Directive was never explicitly drafted with AI in mind. It governs investment services and activities rather than the technologies used to deliver them. Its provisions therefore apply functionally across the regulatory lifecycle—from authorisation and organisational requirements to conduct of business rules, reporting, and supervision. As such, AI falls within MiFID II's scope whenever it mediates a regulated activity, including order execution, portfolio management, investment advice, or the operation of a trading facility.[14]

[12] See, e.g., *Ojanen* [44].
[13] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13.6.2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) [2024] OJ L 1689/1.
[14] See Annex I, Section A, MiFID II.

The following analysis examines how these technology-neutral obligations interact with AI across the principal domains of regulatory relevance under MiFID II. It first considers the horizontal requirements that govern market access, authorisation, and ongoing organisational and reporting obligations, before turning to four key areas of AI application in the financial services ecosystem: (i) consumer-facing applications, (ii) market-facing applications, (iii) firm-internal governance and compliance applications (RegTech), and (iv) supervisory applications (SupTech). Across these areas, the mapping reveals a spectrum of regulatory intensity: while MiFID II's hard law provisions address AI only implicitly—most notably through the rules on algorithmic trading—recent soft law instruments increasingly articulate expectations specific to AI use, whereas in other areas regulatory silence largely prevails.

2.1 Market Access, Authorisation and Other Organisational Requirements

The first point of contact between AI and MiFID II arises at market access during the authorisation process. MiFID II's technology-neutral approach ensures that AI-enabled business models fall squarely within its licensing regime. Under Art. 5 MiFID II, any firm offering investment services or activities, whether delivered by humans or automated systems, must obtain prior authorisation from the relevant national competent authority (NCA). Authorised entities and their permitted activities are recorded in national registers and transmitted to ESMA's central register of investment firms (Art. 5(3)). Authorisation under MiFID II is activity-based: it is triggered by the nature of the service performed rather than the technology employed to perform it. This reflects the enduring regulatory principle of "same activity, same risk, same rules", though the growing use of AI invites reflection on whether automation and data-driven decision-making may, in time, transform the very nature and risk profile of those activities.

A first regulatory interface with AI arises in the programme of operations that firms must submit under Art. 7 MiFID II and Art. 6 of Commission Delegated Regulation (CDR) (EU) 2017/1943.[15] Applicants are required to provide a three-year business plan and detailed information on their organisation, governance, and risk-management arrangements.[16] Although drafted in technology-neutral terms, these provisions are highly relevant for AI-augmented business models. Various regulatory requirements acquire significance for firms deploying financial AI applications in their operations. Firms must, for instance, describe the human and technical resources available, which are likely to include any AI systems integrated into the delivery of investment services as well as other relevant activities.[17] They must also outline any outsourcing arrangements, which may involve, *inter alia*, details of external AI providers, data suppliers, or cloud infrastructures on which operations depend.[18] In addition, firms are required to specify their monitoring and risk control mechanisms, with more demanding obligations applying to those engaging in algorithmic trading within the meaning of Art. 17 MiFID II.[19] Furthermore, applicants are expected to demonstrate how compliance,

[15] Commission Delegated Regulation (EU) 2017/1943 of 14.7.2016 supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards on information and requirements for the authorisation of investment firms [2017] OJ L 276/4. and *Raschner* [45], pp. 107-109.
[16] Art. 6(a), CDR (EU) 2017/1943.
[17] Art. 6(c), CDR (EU) 2017/1943.
[18] Art. 6(e), CDR (EU) 2017/1943. After authorisation, these obligations are effectively extended under Regulation (EU) 2022/2554 (DORA), which imposes detailed, ongoing ICT third-party risk management and register-keeping duties (Arts. 28–30).
[19] Art. 6(g), CDR (EU) 2017/1943.

risk management, and internal audit frameworks capture the use of AI in both client-facing and internal processes,[20] and to show that continuity and resilience planning—including AML/CFT controls—adequately reflects AI-related dependencies and operational risks.[21]

Even without AI-specific provisions, the authorisation process thus functions as a first line of AI governance under MiFID II. It requires firms to demonstrate how technology-related risks are embedded within their overall corporate governance and how associated operational risks are identified, monitored, and mitigated. Firms should expect that NCAs will increasingly scrutinise these aspects of their submissions by seeking clarity on key governance aspects such as model validation, data governance measures, and oversight structures before granting authorisation. Once licensed, investment firms remain subject to MiFID II's organisational requirements (see Art. 16) and to ongoing disclosure and reporting obligations under the broader framework.[22] Supervisory expectations are evolving, with firms now expected to maintain adequate documentation of AI systems used in their operations and to ensure their proper internal review and control.

In sum, MiFID II's authorisation regime demonstrates both the strength and the limits of technological neutrality. AI-powered firms fall clearly within MiFID II's regulatory perimeter, yet the Directive provides little explicit guidance on how AI-specific risks should be assessed or disclosed at the licensing stage. The resulting discretion left to NCAs places growing weight on supervisory interpretation and convergence, particularly through ESMA's coordination and soft-law instruments.

2.2 Consumer-Facing Applications

At the retail level, AI is increasingly embedded in the firm–client interface of investment services. Digital onboarding, robo-advice, portfolio optimisation, and conversational agents now mediate interactions that MiFID II was originally designed to regulate through human advisers. The Directive applies to these innovations through the same functional rules that govern all investment services; it is formally agnostic as to whether services are provided by a person or an algorithm. Yet neutrality does not mean inapplicability, as every AI-assisted interaction must still comply with MiFID II's overarching conduct of business standards—namely, acting honestly, fairly, and professionally in accordance with the client's best interests, and ensuring that all information provided is fair, clear, and not misleading.[23]

The duty to act in the client's best interest under Art. 24(1) binds firms irrespective of delivery mode: automated recommendations, portfolio allocations generated by ML models, or conversational prompts by digital assistants must all serve the client's interest just as traditional advice would. The best interest standard anchors the broader framework of product governance and suitability, ensuring that technological innovation does not dilute investor protection. Arts. 24(2) and 16(3) give concrete expression to these duties through target market and product governance requirements. Accordingly, financial instruments must correspond to the needs, characteristics, and objectives of a defined target market. When AI systems perform profiling or product matching, their logic must reinforce—rather than

[20] Art. 6(h), CDR (EU) 2017/1943.
[21] Arts. 6(i)-(j), CDR (EU) 2017/1943.
[22] See *Raschner* [45], pp. 110-114.
[23] Art. 24 MiFID II.

circumvent—this design. Profiling that extrapolates individual preferences from behavioural data, such as browsing patterns or digital footprints, may create 'hyper-personalised' nudges that subtly steer clients toward products inconsistent with their actual risk appetite. In such cases, the neutrality of the legal text conceals a deeper tension between algorithmic personalisation and investor protection.[24]

Communication duties add a further compliance dimension. Under Arts. 24(3) and 24(5), all information provided to clients must be fair, clear, and not misleading, and presented in a form that clients can understand. This principle applies to every form of communication—whether produced by a human adviser, an automated report generator, or a chatbot integrated into a trading app. AI-generated content must therefore be intelligible both to clients and to internal compliance functions capable of verifying its factual accuracy and compliance with MiFID II's communication standards. GenAI tools, particularly those based on large language models (LLMs) that produce dynamic or personalised content, pose particular challenges, as firms must ensure that such outputs remain explainable, auditable, and consistent with information provided through other delivery channels in order to satisfy MiFID II's standards of clarity and fairness.[25]

The suitability and appropriateness regime under Art. 25 MiFID II and Art. 54 of CDR (EU) 2017/565 equally applies to automated tools. Firms must collect sufficient information about a client's knowledge, financial situation, and risk tolerance, and ensure that any recommendation or decision is suitable. The 2023 ESMA *Guidelines on Suitability* confirm that these duties apply without attenuation to hybrid or fully automated systems. While AI tools may help firms collect and analyse client information or assess risks more efficiently, they do not lessen firms' responsibility for the final suitability outcome.[26] A related interpretative issue concerns the legal qualification of the "advisor". Art. 25(1) attributes advisory competence to natural persons deemed qualified to provide investment advice. The 2023 ESMA *Briefing on the Definition of Investment Advice* reiterates that machines cannot hold this legal status. It also reaffirms that responsibility remains with the authorised firm and its human management. Hence, automation may assist, but it cannot displace human accountability or legal liability.[27] The 2024 ESMA *Public Statement on AI in Retail Investment Services* further clarifies that clients should be informed whenever they interact with an AI system[28]—a principle that echoes Art. 22 General Data Protection Regulation (GDPR)[29] on automated decision-making and, for high-risk AI systems, Art. 86 of the AI Act, which similarly requires transparency and meaningful disclosure to affected persons.

The growing reliance on AI in client-facing services also generates risks that MiFID II's existing provisions only partially address. Profiling and behavioural targeting—techniques that often remain opaque to clients—may lead to hyper-nudging, price discrimination, or biased filtering, each of which can distort suitability assessments and compromise the duty to treat clients fairly. Chatbots and virtual assistants can miscommunicate, omit critical warnings, or fail to escalate complex cases for human review. AI-generated communications—such as personalised reports or educational nudges—also introduce risks of persuasive framing,

[24] *Buczynski et al.* [13], p. 7.
[25] *Id.*; see also *Gehrmann et al.* [35].
[26] ESMA [31].
[27] ESMA [30].
[28] ESMA [29].
[29] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27.4.2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC [2016] OJ L 119/1.

factual hallucinations, or inconsistencies across delivery channels. All such phenomena challenge the Directive's expectation that the information provided to clients must be correct, comprehensible, and accurate. In response, supervisory authorities have begun translating general principles into expectations for AI governance. The 2024 ESMA *Public Statement on AI in Retail Investment Services* identifies at least six key supervisory priorities, including (i) ultimate board responsibility for AI tools; (ii) robust risk-management and testing frameworks; (iii) client transparency and disclosure of AI use; (iv) assurance of data quality and bias control; (v) meaningful human oversight; and (vi) adequate staff training to monitor and, where necessary, intervene in AI processes.[30] Although non-binding, these initiatives fill interpretative gaps left by technology-neutral regulation and signal a shift toward more concrete, risk-sensitive expectations for AI oversight.

2.3 Market-Facing Applications

The most explicit interaction between MiFID II and automated technologies arises in algorithmic trading and its subfield of high-frequency trading (HFT).[31] In addressing algorithmic trading governance, Art. 17 MiFID II and the detailed standards in CDR (EU) 2017/589 (RTS 6)[32] constitute the most technologically specific component of the MiFID II framework. Introduced in response to the 2010 "flash crash", these provisions aim to ensure that trading technologies operate in a manner consistent with fair and orderly markets and with the requirements of the Market Abuse Regulation (MAR).[33]

Under Art. 17(1), firms engaging in algorithmic trading must maintain effective systems and risk controls to ensure that their trading engines are resilient, operate within defined parameters, and comply with MAR. They must also have sound business continuity plans, pre-deployment testing, and post-deployment oversight. In principle, these obligations apply equally to deterministic algorithms and to adaptive, self-learning systems. Yet aligning an ML-powered trading system that updates its behaviour in response to market experience with MiFID II's *ex ante* governance requirements presents qualitatively different challenges from those posed by more rudimentary systems. Safeguarding against the risk that ML-based systems generate disorderly trades or exploit emergent market patterns requires a level of behavioural control that exceeds what the original legislation was designed to address.[34]

Transparency obligations under Art. 17(2) require firms to notify their NCA when employing algorithmic trading and, upon request, to provide detailed descriptions of strategies, parameters, risk controls, and testing arrangements. High-frequency traders must additionally keep time-sequenced records enabling firms and supervisors to reconstruct trading behaviour and verify compliance. In practice, such transparency is difficult to achieve where trading systems are opaque or adaptive, particularly when powered by ML methods. The "nature of the algorithmic strategy" may be unintelligible even to its developers, rendering disclosure

[30] See *id.*

[31] This discussion is necessarily non-exhaustive. Other relevant applications—such as AI adoption in regulated markets, direct electronic access, and order-routing—lie beyond the present scope for reasons of space.

[32] Commission Delegated Regulation (EU) 2017/589 of 19.7.2016 supplementing Directive 2014/65/EU of the European Parliament and of the Council with regard to regulatory technical standards specifying the organisational requirements of investment firms engaged in algorithmic trading [2017] OJ L 87/417.

[33] Regulation (EU) No 596/2014 of the European Parliament and of the Council of 16.4.2014 on market abuse (market abuse regulation) and repealing Directive 2003/6/EC of the European Parliament and of the Council and Commission Directives 2003/124/EC, 2003/125/EC and 2004/72/EC [2014] OJ L 173/1 and Recitals (61) (64) MIFID II.

[34] *Azzutti et al.* [9], pp. 123-126.

partial and compliance with the spirit of Art. 17 largely aspirational. RTS 6 complements these duties with dense organisational and procedural safeguards. It requires behavioural and conformance testing before deployment, re-testing following substantial updates, and the establishment of "hard limits" defining the instruments, venues, and price thresholds accessible to each algorithm. Firms must also implement pre-trade, real-time, and post-trade controls and maintain surveillance mechanisms capable of detecting abusive or anomalous conduct. Taken together, these provisions amount to a quasi-lifecycle discipline for AI trading. Yet their logic presupposes static, rule-defined architectures whose behaviour can be validated *ex ante*. Adaptive ML systems, however, unsettle that assumption, as models evolve through interaction with live data. Their market behaviour may diverge from pre-tested parameters, leaving firms and supervisors uncertain as to when re-validation obligations are triggered and complicating the definition of effective constraints where models autonomously adjust trading strategies within boundaries that are difficult to predefine.[35]

The regulatory gap is particularly evident in the area of market conduct. MAR Art. 12 defines market manipulation through specific behavioural patterns and market effects, but its construction implicitly assumes human agency and rational motivation. This framework becomes problematic when applied to autonomous trading systems capable of developing harmful or manipulative strategies without direct human direction.[36] RTS 6 Art. 13 requires firms to operate surveillance systems capable of detecting abusive patterns, while MAR Art. 16 imposes an obligation to report suspicious transactions or orders to the competent authority or trading venue. Yet many of these monitoring tools increasingly rely on ML methods, thereby compounding the transparency problem: supervising AI with AI introduces an additional layer of opacity and explainability risk.[37] In practice, MiFID II's algorithmic trading governance framework remains both agnostic to the specific AI technologies employed and outcome-based. By focusing on observable conduct rather than the underlying processes, its functional breadth is increasingly strained by adaptive and data-driven systems—adequate for rule-based automation, less so for systems whose internal logic cannot be fully documented or audited.

While MiFID II's framework on algorithmic trading provides an impressive degree of functional coverage for AI applications, it only offers a partial response to the distinctive risks that accompany more advanced trading systems. First, sophisticated learning architectures—such as reinforcement learning-based trading agents—may autonomously develop manipulative or collusive trading patterns, including spoofing or quote-stuffing, without human intent. Although MAR Art. 12 does not require proof of intent, its underlying logic presumes human decision-making, which sits uneasily with autonomous behaviour and raises questions about liability attribution. Second, the opacity of black-box models undermines both regulatory disclosure and internal compliance. The obligation to describe the nature and certain other details of algorithmic strategies under Art. 17(2) assumes a level of transparency or explainability that many AI systems cannot deliver. This lack of insight, in turn, can impair firms' ability to identify and substantiate suspicious transactions, escalate potential misconduct internally, and provide meaningful information to supervisors. Third, phenomena such as model or data drift, as well as other emergent behaviours inherent in adaptive AI, challenge the stability assumed by RTS 6. In such cases, incremental changes may not qualify as

[35] *Azzutti* [6], pp. 64 and 75.
[36] *Azzutti et al.* [9]; *Azzutti* [6].
[37] *Azzutti* [7].

"substantial updates", even though they materially alter trading behaviour. Finally, informational asymmetries may further widen as supervisory authorities confront systems that neither firms nor regulators can fully interpret or audit.[38]

Viewed through the prism of algorithmic trading, MiFID II's technology-neutral framework shows a dual character. On the one hand, it is *functionally enabling*, in that it is sufficiently broad to capture AI applications in automated trading. On the other hand, it is *normatively limited*, as it rests on assumptions of human agency and oversight, as well as predictable system behaviour, that more advanced generations of AI trading may no longer satisfy. While certain risks associated with automated technologies are already embedded within the existing governance framework, the latter provides only partial tools for addressing key techno-methodical characteristics of ML-powered systems—particularly their opacity, adaptivity, and potential for emergent behaviour. This, in turn, exposes the limits of a static hard-law approach to evolving AI risks.

2.4 Firm-Internal Governance and Compliance Applications (RegTech)

AI is increasingly embedded in the internal governance and compliance functions that MiFID II regulates through its organisational requirements. These duties, primarily set out in Art. 16 MiFID II and related provisions, focus on sound governance, effective internal control, and reliable record-keeping, without prescribing the technologies by which these aims must be achieved. This non-prescriptive yet enabling design has provided fertile ground for the rise of Regulatory Technology (RegTech)—a diverse field of tools that automate compliance, risk assessment, and reporting across the governance, risk, and compliance (GRC) spectrum.[39]

Within the organisational requirements under Art. 16, firms increasingly rely on AI systems to identify compliance risks, flag breaches, monitor algorithmic trading, and support broader risk management and internal audit processes. Such tools can detect anomalies or control breaches in real time, maintain digital audit trails, and assist human experts in exercising supervisory and decision-making tasks. They enhance efficiency, consistency, and auditability, but also create new dependencies on data quality, model integrity, cybersecurity resilience, and third-party providers—factors that MiFID II does not always explicitly address. Further intersections arise under Art. 24, which governs client information and communications. AI-based language processing and document generation tools help ensure disclosures are accurate and consistent with regulatory templates. Such systems can also monitor whether digital communications—ranging from website content to automated chat or messaging systems—remain within the boundaries of acceptable conduct. In principle, AI-driven RegTech solutions are intended to enhance compliance by standardising outputs and reducing human error; in practice, however, their reliability depends critically on the transparency, validation, and governance of the underlying data, models, and technical infrastructure. Comparable uses arise across other conduct of business obligations. Under Art. 25, which establishes the suitability and appropriateness regime, AI systems are increasingly used to perform dynamic client-risk profiling, flag mismatches between products and investor characteristics, and guide decision-making towards compliant recommendations. When responsibly designed, these

[38] *Azzutti et al.* [9]; *Azzutti* [7]; see also *Annunziata* [3].
[39] See The RegTech Association and Deloitte [46].

tools can strengthen investor protection by improving the accuracy and consistency of suitability checks. Yet excessive reliance on opaque or unvalidated systems risks turning compliance into a formalistic exercise detached from truly accountable human oversight. A similar pattern emerges in the best-execution framework of Art. 27, where AI is used to analyse transaction data, benchmark execution quality, and identify routing or pricing anomalies. While such systems increase analytical depth and speed, they rely on continuous calibration, high-quality data, and access to sensitive information, including personal data, thereby raising concerns about data governance, confidentiality, and vendor oversight.

Beyond these core obligations, RegTech adoption extends to functions that the Directive does not directly regulate but that increasingly affect how firms achieve compliance—most notably, automated horizon scanning and cross-regime monitoring. AI tools—including, increasingly, GenAI—are now deployed to identify, map, and reconcile overlapping legal obligations under MiFID II, the Digital Operational Resilience Act (DORA)[40], the GDPR, AML rules, and sustainability-related disclosure, among others. This regulatory overlap exposes a deeper structural tension, insofar as different legal regimes may pursue conflicting policy objectives. For instance, while AML laws promote comprehensive data collection and retention to detect illicit activity, the GDPR imposes duties of data minimisation and purpose limitation. Reconciling these objectives requires firms to balance surveillance and privacy requirements. AI-based compliance tools operating amid these competing imperatives may comply with one regime while undermining another. Beyond these cross-regime tensions, an additional layer of operational and supervisory challenges has emerged within firms' own compliance infrastructures. AI adoption remains uneven across firms and Member States, reflecting the absence of harmonised expectations regarding the use of AI in regulatory compliance. Heavy reliance on external vendors heightens outsourcing risks under Art. 16(5), particularly where firms lack visibility into proprietary models or data processing practices. Automation bias may further erode human judgment, as compliance staff over-rely on system outputs without adequate scrutiny.[41]

Above all, MiFID II and its Level 2 measures define regulatory outcomes but not the technical means of achieving them. While they require effective internal controls, they offer little guidance on how AI-based compliance systems should be designed, validated, or audited. Recent supervisory communications have at least begun to fill this gap by emphasising management accountability, model testing and validation, data-quality assurance, and explainability, but these remain principles rather than binding standards.[42] As a result, the governance of AI-driven RegTech currently occupies a grey zone between innovation and regulatory certainty. Firms can demonstrate formal compliance with MiFID II's requirements, yet the underlying technological processes often remain opaque or not fully verified. This ambiguity sustains MiFID II's flexibility but leaves firms and supervisors uncertain about how the safe and trustworthy use of AI in compliance should be operationalised.

[40] Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14.12.2022 on digital operational resilience for the financial sector and amending Regulations (EC) No 1060/2009, (EU) No 648/2012, (EU) No 600/2014, (EU) No 909/2014 and (EU) 2016/1011 [2022] OJ L 333/1.

[41] Cf. EBA [20], pp. 14-18, for a discussion of these challenges in the context of RegTech adoption in AML and financial crime prevention.

[42] See, e.g., ESMA [29]; EBA [21].

2.5 Supervisory Applications (SupTech)

AI is not confined to the private sphere of financial firms. Supervisory authorities themselves increasingly rely on AI to manage the scale, speed, and complexity of supervisory tasks arising from the post-MiFID II regulatory framework. While MiFID II does not seek to regulate supervisory technologies as such, the use of AI within so-called Supervisory Technology (SupTech) has become an important enabler of its underlying objectives. Rather than constituting a distinct regulatory domain, SupTech has emerged as a strategic response to the data-intensive and analytically demanding nature of EU financial supervision.

MiFID II and its companion regimes—most notably the Markets in Financial Instruments Regulation (MiFIR)[43], but also MAR, the European Market Infrastructure Regulation (EMIR)[44], the Securities Financing Transactions Regulation (SFTR)[45], and, more recently, DORA—collectively generate vast volumes of structured and unstructured supervisory data. The resulting transaction reports, product governance disclosures, market-abuse alerts, and ICT incident notifications pose a significant supervisory challenge while simultaneously enabling more data-driven oversight. NCAs increasingly experiment with ML tools to detect anomalies, map trading networks, and prioritise investigations, while ESMA plays a coordinating role in promoting supervisory convergence and data-driven innovation. These developments promise greater efficiency and responsiveness but also raise fundamental questions of legality and accountability.

As SupTech lies outside MiFID II's direct remit, its normative foundation instead rests on EU administrative law, including the principle of good administration. This principle embodies duties of due care and due process—i.e. fairness, timeliness, impartiality, the right to be heard, and reason-giving—which continue to bind authorities even when algorithms inform supervisory or enforcement decisions. In this context, explainability, accountability, and auditability become essential procedural guarantees for legitimate supervisory action. For instance, an AI-generated alert may legitimately inform or prioritise supervisory inquiries, but ensuing decisions must remain human, reasoned, and reviewable. Excessive automation risks undermining the discretion and proportionality that underpin good administration.[46] Beyond procedural legality, algorithmic supervision raises questions of institutional capacity. Effective SupTech deployment presupposes not only adequate technical infrastructure but also expertise in data science, model validation, and ethical governance, as well as other key technical competencies and organisational capabilities. Supervisory authorities must therefore establish safeguards (specifically in relation to testing, documentation, and oversight) that are comparable to—and arguably superior to—those expected of the firms they regulate. In this reflexive sense, supervisory accountability should mirror, if not exceed, the accountability demanded of supervised entities.[47]

Overall, SupTech applications epitomise the dual legacy of MiFID II's technology-neutral architecture. The Directive's emphasis on transparency and supervisory data availability has

[43] Regulation (EU) No 600/2014 of the European Parliament and of the Council of 15.5.2014 on markets in financial instruments and amending Regulation (EU) No 648/2012 [2014] OJ L 173/84.
[44] Regulation (EU) No 648/2012 of the European Parliament and of the Council of 4.7.2012 on OTC derivatives, central counterparties and trade repositories [2012] OJ L 201/1.
[45] Regulation (EU) 2015/2365 of the European Parliament and of the Council of 25.11.2015 on transparency of securities financing transactions and of reuse and amending Regulation (EU) No 648/2012 [2015] OJ L 337/1.
[46] *Azzutti et al.* [8].
[47] See, e.g., *Barasa et al.* [10].

compelled regulators to innovate technologically, even though its core objectives were never designed to prescribe whether, or how, AI and related technologies should be adopted by financial supervisors. The same tension visible at the firm level—i.e. functional enablement coupled with normative silence—reappears within public administration itself. While efficiency gains are tangible, the legitimacy of algorithmic supervision ultimately depends on preserving the procedural and substantive standards that underpin the rule of law in EU public administration.

3. From Neutrality to Nuance: Risk, Proportionality, and the Future of AI Governance in Investment Services

The analysis above shows that MiFID II intersects with AI applications across all key domains of regulatory relevance—ranging from firm-consumer interactions and firm-market activities to firm-internal governance and compliance processes, and finally regulator-firm supervisory relationships. Across these domains, MiFID II operates through a framework that is functionally enabling yet (largely) normatively silent. Rooted in the principle of technological neutrality, the Directive regulates functions rather than technologies. This design has long been praised for "future-proofing" legislation by ensuring that regulation follows activity rather than technological form.48 Yet the diversity of financial AI applications now challenges that very assumption. As Section 2 illustrated, algorithmic trading is governed by detailed hard law provisions, consumer-facing AI by emerging soft law guidance, and RegTech (as well as SupTech) by only minimal normative anchoring. The result is a patchwork of regulatory intensity that reflects MiFID II's flexibility but also exposes its limits in addressing technology-specific risks. Across and within these application domains, treating heterogeneous AI systems as functionally equivalent conceals material differences in their risk profiles—ranging, inter alia, from investor protection and market integrity to operational resilience—thereby testing the adequacy of technological neutrality as a stand-alone regulatory principle. The remainder of this section examines that tension, first by analysing the limits of neutrality and then by outlining how a more proportionate, risk-based approach could complement, refine, and, where necessary, adapt the existing regulatory framework.

3.1 Financial AI Governance and the Limits of Technological Neutrality

MiFID II's principle of technological neutrality was likely conceived for a financial system in which AI and related technologies augmented human judgement within existing service workflows, rather than transforming the logic through which financial services and ancillary tasks are performed. Its functional logic presumes that equivalent activities can be governed by the same regulatory obligations, regardless of the technological means through which they are performed.[49] That premise, however, begins to falter once those means take the form of

[48] For a critical account, see *Ojanen* [44].

[49] For recent Court of Justice of the European Union references to technological neutrality in the interpretation of EU law, see Case C-135/23 *GEMA*, para. 37, EU:C:2024:526; Case C-433/20 *Austro-Mechana*, para. 27, EU:C:2022:217; Case C-515/19 *Eutelsat*, para. 48, EU:C:2021:273.

autonomous systems capable of self-learning, adaptation, and action beyond meaningful human oversight.[50]

As discussed in Section 2, this tension is already visible across domains. In market-facing contexts, while the hard law framework of Art. 17 and RTS 6 remains broadly adequate for deterministic algorithmic trading, it appears increasingly misaligned with adaptive AI systems that evolve dynamically through interaction with market and other data. In consumer-facing applications, neutrality ensures that AI-based or 'robo-'advice is subject to the same legal standards as human advice; yet it leaves unaddressed the lifecycle risks specific to more advanced AI systems—such as, *inter alia*, explainability, profiling bias, and behavioural steering—that ultimately determine whether such advice is trustworthy and compliant in practice. ESMA's recent soft law guidance in this area shows that, while conduct of business regime remains formally technology-neutral, *de facto* supervisory expectations are becoming increasingly rigorous for firms deploying AI at the client interface. RegTech applications further expose this asymmetry. While MiFID II defines the regulatory outcomes firms must achieve, it offers little guidance on how those outcomes are to be generated or verified when the underlying processes operate beyond meaningful human oversight. This open formulation is deliberate: by focusing on outcomes rather than prescribed processes, MiFID II grants firms a certain degree of flexibility in determining the most efficient means of compliance. Yet that flexibility becomes problematic where the compliance processes themselves—such as, *inter alia*, AI model design, validation, or monitoring—introduce new layers of uncertainty that the law does not fully and explicitly anticipate. A similar logic applies to SupTech, where supervisory authorities confront parallel challenges of explainability and control. The key difference, however, lies in the legal and institutional framework governing AI adoption within supervisory bodies' public oversight functions, as opposed to its use within firms' internal compliance systems.

Across these different domains, a consistent pattern emerges. While MiFID II's technology-neutral design ensures formal inclusion of AI within existing legal categories of financial services and activities, it fails to align substantively with the distinct risk profiles that characterise the diverse spectrum of financial AI applications. Many such applications fall within the Directive's perimeter, yet their heterogeneous dimensions of risk—spanning, *inter alia*, consumer and investor protection, market safety and integrity, operational resilience, and data governance—are only tangentially addressed. Equal treatment of *unequal* technological applications thus produces a form of regulatory asymmetry, whereby the same legal rules govern systems with markedly different technical architectures, capabilities, and levels of risk materiality. This lack of differentiation leaves both firms and supervisors navigating an increasingly under-specified legal framework, resulting in heightened compliance challenges for firms and divergent supervisory interpretations and uneven oversight across Member States. This asymmetry also reveals a deeper conceptual tension between technological neutrality and regulatory proportionality. While the former ensures formal equality by applying the same rules to all technologies, the latter requires that regulatory intensity reflects the specific risks and impacts of a given activity. MiFID II already incorporates a limited form of risk calibration—most clearly in Art. 17, which prescribes granular governance standards for algorithmic trading. There, legislators recognised that automation warranted distinct regulatory treatment, though that calibration now struggles to address the complexity of more

[50] See, e.g., *Azzutti* [5].

advanced systems.[51] Yet this insight remains confined to a single domain. Elsewhere, neutrality operates as a blanket assumption, treating applications like AI-based advisory systems, portfolio optimisation tools, and RegTech solutions as if they carried equivalent risks as traditional processes, based on the assumption that existing conduct of business and organisational duties are sufficient to manage them. This contrast underscores the limits of MiFID II's neutrality approach. Where proportionality has been applied, regulatory expectations are, at least in principle, clearer and more enforceable. Where neutrality prevails, by contrast, they remain broad, fragmented, and largely reactive, articulated only through *ex post* guidance.

In light of the foregoing, technological neutrality thus functions as a principle of activity inclusion but not of risk calibration. While it guarantees that AI-based financial services fall within the scope of MiFID II, it leaves their differentiated risks largely unaccounted for. As a result, what was once praised for future-proofing financial legislation now risks generating conceptual and normative gaps. As AI systems are deployed across increasingly diverse business functions and reach new levels of technical sophistication and complexity—from adaptive trading algorithms to generative advisory agents—the assumption that uniform functional rules can adequately govern such heterogeneous technological applications becomes increasingly difficult to sustain. The next step, therefore, is to consider how technological neutrality might be complemented by a risk-based governance approach that preserves MiFID II's flexibility while enabling proportionate responses to the risks posed by specific AI applications.

3.2 Towards a Risk-Based Taxonomy for Financial AI Applications

The limitations of technological neutrality do not call for a wholesale replacement of MiFID II but for its gradual evolution toward a risk-based model of technological governance better suited to financial AI. Such an approach would preserve the Directive's functional flexibility—its capacity to accommodate technological change within stable legal categories of regulated activity—while introducing appropriate sensitivity to the heterogeneous and evolving risk profiles of AI systems in specific contexts of application. The doctrinal foundation for this shift lies in the EU law principle of proportionality, which requires that legal obligations be calibrated to the magnitude and systemic significance of the risks at stake.[52] As previously noted, algorithmic trading under Art. 17 illustrates an initial attempt at risk calibration within the existing framework. Legislators recognised automation in financial trading as a risk amplifier and imposed tailored governance requirements. Yet even within this domain, the framework increasingly struggles to address the additional market and technological complexity introduced by the latest generations of AI.[53] The need, therefore, is not only to extend the principle of proportionality horizontally across distinct domains of AI use but also to refine it vertically within existing regimes, thus ensuring that regulatory intensity evolves in line with the technological risks and systemic relevance of specific applications.

[51] *Id.*

[52] For a critical discussion on the role of proportionality in the evolution of the EU approach in regulating finance, particularly the banking sector, see *Buttigieg/Armeni Cauchi* [14].

[53] *Azzutti* [5].

Across EU digital finance governance, risk-based regulation has emerged as a clear pattern. Most notably, the AI Act establishes a horizontal framework that classifies AI systems into four tiers of risk—unacceptable, high, limited, and minimal—and attaches proportionate obligations across the AI lifecycle, including, *inter alia*, requirements relating to risk management, data governance, technical documentation, transparency, and human oversight. Along the AI value chain, these obligations fall most heavily on providers and deployers of high-risk AI systems, while applying with more limited—and largely transparency-based—intensity to importers, distributors, and users. Interestingly enough, the only explicitly financial use cases designated as high-risk under the Act concern credit scoring and access to essential financial services, leaving other applications—such as trading, advisory, or compliance tools—largely outside its stricter tier.[54] A similar logic of proportionality operates in the GDPR, where compliance obligations scale with the risks that processing poses to individuals' rights and freedoms,[55] and in DORA, which differentiates duties for ICT providers according to their criticality and systemic interdependence.[56] Meanwhile, recent ESMA guidance on AI in investment services—though soft law—points toward an emerging model of AI-specific governance grounded in adequate risk management and control, meaningful human oversight, and appropriate AI literacy within firms.[57] Taken together, these developments indicate a broader regulatory shift from technological neutrality toward risk-calibrated proportionality. Yet despite this convergence, the financial sector still lacks a coherent taxonomy for translating that proportional logic into its own sectoral framework.

Beyond investment services, several EU regulators have already begun operationalising risk-based taxonomies for AI within their respective sectors. The European Medicines Agency (EMA) distinguishes between patient safety risk and regulatory impact risk across the medicinal product lifecycle.[58] The Medical Device Coordination Group (MDCG) has refined software classification under the MDR/IVDR to capture AI-enabled medical devices through graded risk classes.[59] The European Union Aviation Safety Agency (EASA) calibrates assurance requirements according to the degree of human–AI interaction in specific applications.[60] In parallel, the EBA and the European Insurance and Occupational Pensions Authority (EIOPA) have advanced sector-specific approaches—respectively, a harmonised operational-risk event taxonomy encompassing AI incidents,[61] and a proportionate governance framework that scale controls to the materiality of AI use.[62] Altogether, these initiatives illustrate a pan-sectoral regulatory movement toward risk-based taxonomies that complement the horizontal logic of the AI Act and demonstrate the feasibility of transposing that logic into financial services.

A coherent risk-based framework for financial AI should ideally combine horizontal baseline requirements with proportional, domain-sensitive obligations.[63] The objective is not to multiply regulated categories of activity, but to align the intensity of legal obligations with the

[54] Annex III, AI Act. In principle, the obligations applicable to general-purpose AI (GPAI) systems under the AI Act also extend to the financial sector, particularly where firms act as providers of such tools, including—where relevant—those deemed to pose 'systemic risk'. However, this paper does not examine GPAI-specific requirements in detail, but instead focuses on the sectoral governance of AI within the MiFID II framework.
[55] For a critical account, see *Della Corte* [16].
[56] *See Wittlin et al.* [48]; *Clausmeier* [15].
[57] ESMA [29].
[58] EMA [24].
[59] MDCG [40].
[60] EASA [17].
[61] EBA [19].
[62] EIOPA [23].
[63] Cf. *Azzutti* [5], pp. 28-29, and sources cited therein.

nature and magnitude of AI-related risks as they materialise in specific application contexts. Horizontally, all AI systems operating within, or functionally interacting with, the MiFID II perimeter—whether client-facing, market-facing, or used for RegTech purposes—should comply with a core set of governance principles, for instance on transparency, accountability, explainability, cybersecurity, and auditability. These baseline duties would apply uniformly, ensuring a consistent standard of AI trustworthiness throughout the financial ecosystem. Proportional differentiation would then operate within this common baseline, scaling regulatory expectations to the specific risks embedded in each application. The intensity of those risks depends not only on the system's function but also on its techno-methodical configuration—its design logic, operational capabilities and limits, and the organisational or infrastructural context in which it is deployed. In this sense, horizontal classification of AI by functional domain helps to clarify both the *nature* of risks likely to arise and their primary *loci* of exposure: i.e. consumers and investors in client-facing contexts; market integrity and stability in market-facing applications; internal governance and operational resilience in RegTech; and institutional legitimacy and accountability in SupTech.

Building on this foundation, the proportional layer of governance requirements would thus capture the magnitude and dynamics of specific risk through the combination of two further analytical dimensions:

(i) The *source* of risk identifies where potential harm originates, as determined by the techno-methodical complexity of the AI system and its broader value chain. This dimension measures the difficulty of effective governance, tracing risk across two primary layers: 'intrinsic complexity' (e.g., degree of autonomy, adaptivity, and model opacity) and 'structural complexity' (e.g., data dependencies, integration into critical infrastructures, and reliance on third-party or cross-border providers). These determinants reflects both the intra- and inter-organisational nature of the AI lifecycle, acknowledging that the source of risk may often reside in emergent behaviours of system architectures and the opacity of supply chains.

(ii) The *severity* of risk measures the magnitude, diffusion, and reversibility of potential harm across the financial system. This dimension evaluates the potential depth of impact—ranging from isolated operational incidents or impacts on specific client groups to cross-market or systemic disruptions. This dimension depends largely on the system's capabilities (e.g., predictive power, speed of execution, and cognitive reach), as well as its action space (i.e. the range of decisions the system is permitted to take within its operational environment). Hence, severity is a function of the propagation potential of the harm cause by AI wrongdoing.

These dimensions give analytical content to the notion of 'risk levels' and enable regulators to differentiate both across functions and among varying degrees and types of exposure. They also provide a consistent framework for firms to assess, document, and communicate their use of AI systems. At the foundation lies a simple yet essential requirement: each firm should maintain an *AI inventory* identifying all regulatory-relevant systems in use, their functional domains, and their assessed risk profiles.[64] Without such visibility, supervisors cannot

[64] For a high-level introduction to the role of AI inventories in governance, see *Wendt* [47], p. 104. The notion of maintaining system inventories already exists under MiFID II and RTS 6, which require firms engaged in algorithmic trading to document and record their trading algorithms. A similar rationale underpins the EU AI Act's public database for high-risk AI systems; *see Buczynski et al.* [13], p. 10. The present proposal extends this principle to firm-level governance, requiring financial institutions to maintain internal AI inventories of regulatory-relevant systems—proportionate to risk and materiality—for both compliance and supervisory purposes.

effectively map, compare, or monitor the evolving landscape of financial AI.[65] The AI inventory thus provides the factual baseline for applying proportionality in practice, anchoring both firm-level governance and supervisory oversight to the actual scope and risk profile of technological deployment.[66] Table 1 below provides a schematic visualisation of the analytical dimensions that structure a risk-based assessment of financial AI systems.

Table 1. Analytical Framework for a Risk-Based Assessment of Financial AI

Dimension	Purpose and Focus	Risk Classification Scale (Example)
Nature of Risk	Identifies the *locus of exposure* and the specific legal interest.	Micro (Consumer) → Meso (Firm) → Macro (Market)
Source of Risk	Pinpoints the *techno-methodical origin* within an AI system's design, operational environment, or value chain.[67]	Low → Medium → High Complexity
Severity of risk	Measures the *magnitude* and *propagation potential* of harm across the system, as determined by system capabilities and operational context.[68]	Minimal → Low → Medium → High → Unacceptable

The proposed taxonomy moves beyond simple categorisation to embed a dual logic of *AI safety* and *AI security* directly within the scope and objectives of financial regulation. By decoupling risk intensity from specific technological labels, it is designed to future-proof the governance of financial AI applications. Externally, the framework contributes to strengthen market and consumer protection by enabling supervisors to better identify, prioritise, and mitigate risks that could compromise investor interests, market integrity, or even financial stability. Internally, it promotes the resilience of firms' technological infrastructures by ensuring that AI systems are developed and operated in ways that are robust, transparent, auditable, controllable, and cyber-secure. Ultimately, this approach displaces 'one-size-fits-all' regulation. By structuring oversight around risk intensity rather than specific technologies, it also offers tangible advantages for firms, including clearer compliance benchmarks, more predictable supervisory expectations, and a more proportionate allocation of oversight resources.

There are, of course, multiple ways to design a risk-based governance framework for financial AI, and recent scholarship has already begun to explore graduated or proportional models of AI governance.[69] Nevertheless, any attempt to introduce a finely calibrated risk-based taxonomy inevitably encounters the 'simplicity–complexity' trade-off in regulation.[70]

[65] See IOSCO [36], p. 33; OECD [43], p. 13, both noting persistent "knowledge gaps" among regulators regarding firms' AI use and controls. A similar concern underlies ESMA's 2025 survey on AI in investment services, launched to collect industry data and enhance supervisory visibility over emerging AI deployment.

[66] This proposal is increasingly reflected in regulatory developments, particularly the 2025 consultation by the Monetary Authority of Singapore (MAS) on draft Guidelines for AI Risk Management, which invites financial institutions to maintain a comprehensive AI inventory. See MAS [18].

[67] Relevant criteria may include, *inter alia*, degree of autonomy or adaptivity; data dependency and sensitivity; model transparency; integration into critical infrastructures; reliance on third-party or cross-border providers.

[68] Relevant criteria may include, *inter alia*, scale of affected clients or entities; reversibility and duration of effects; interconnection with other systems or markets.

[69] See, e.g., *Azzutti* [5], however exclusively focusing on algorithmic trading and HFT domains.

[70] See *Kaplow* [37]; *Epstein* [25].

Increasing granularity enhances precision and proportionality, yet it may also magnify interpretive and operational burdens.[71] Overly detailed taxonomies risk fragmenting supervision or generating disproportionate compliance costs, while excessive simplification can obscure material differences in the sources, manifestations, and transmission of risk. Complexity, however, is not inherently pathological; it tends to reflect the structural features of contemporary legal systems tasked with governing intricate socio-technical realities.[72] The central challenge is therefore not to eliminate complexity but to govern it deliberately[73]—to achieve *engineered simplicity*: a regulatory framework that remains intelligible, predictable, and enforceable while retaining sufficient nuance to capture the diversity of financial AI systems.[74] Whether expressed through rules or principles, this balance depends on norms that are both clear and adaptable under conditions of uncertainty, supported by mechanisms for iterative refinement and evidence-based adjustment.[75]

An improved risk-based framework would allow MiFID II to retain its functional flexibility while addressing the normative silence identified earlier in this paper. By embedding proportionality into its operational logic, MiFID II could evolve from a model of formal neutrality to one of structured risk-sensitivity, in which regulatory obligations scale dynamically with the risks inherent in specific AI systems. The proposed taxonomy would serve both interpretative and regulatory purposes. Interpretatively, it reframes MiFID II obligations through a renewed proportionality lens, enabling supervisors and firms to align regulatory expectations, compliance mechanisms, and oversight practices with the actual characteristics and unique risks of different applications. Regulatively, it may offer a structured template to inform future RTS, supervisory guidance, or legislative reform. Its role is thus both conceptual—clarifying how proportionality applies to AI—and practical—embedding that principle in the operational logic of risk governance. While inspired by the AI Act's tiered classification, the proposed taxonomy remains sector-specific in scope and purpose. The AI Act primarily safeguards EU citizens' fundamental rights and safety, whereas financial regulation pursues investor protection, market quality and integrity, and systemic stability. These normative objectives overlap only partially. Adapting the AI Act's structural logic to the financial context therefore promotes coherence across the EU's approach to AI governance without conflating distinct regulatory rationales. In this sense, a risk-based taxonomy for financial AI applications would complement the horizontal AI regime by embedding proportionality within the sector's own 'risk ecology'—the interdependent configuration of technological, market, and organisational risks that shape financial AI.

Having outlined the conceptual basis, the next question concerns implementation. Translating this conceptual foundation into practice need not require a legislative overhaul. A progressive implementation strategy can act as a mechanism of complexity management, introducing differentiation gradually and allowing supervisory capacity and empirical knowledge to evolve alongside regulatory sophistication. In the short term, soft law guidance from EU authorities could clarify risk-assessment criteria for AI systems across functional domains, thereby fostering supervisory convergence. In the medium term, Regulatory and Implementing Technical Standards (RTS/ITS) could formalise proportionate documentation, testing, and disclosure duties calibrated to technology-specific risk. Over the longer term, if these measures

[71] See BIS [11].
[72] Cf. *Appermont* [3].
[73] *Black/Baldwin* [12].
[74] Cf. *Aikman et al.* [2].
[75] Cf. *European Commission* [33].

prove insufficient, a future MiFID III reform could codify risk-based proportionality within the legislative framework.[76] While regulatory failure in finance can entail severe societal costs, a staged approach would respect the deliberative nature of EU law-making while ensuring that regulatory evolution remains coherent, iterative, and empirically grounded.

In this context, regulatory sandboxes could play an important role in advancing and operationalising an adaptive, risk-based framework for financial AI. While several Member States have established national innovation hubs and sandbox initiatives, regulatory experimentation and learning remains largely decentralised, with no fully institutionalised EU-level mechanism for systematic coordination or cross-border knowledge transfer.[77] As a result, valuable supervisory insights tend to remain fragmented. Comparative experience nevertheless illustrates the potential of sandbox-based approaches. In the United Kingdom, for example, the Financial Conduct Authority has developed a comparatively mature sandbox ecosystem within a broader principle-based approach to digital innovation, of which the Consumer Duty framework is a prominent illustration. More than mere environments for technical testing, sandboxes can function as controlled settings in which regulators and market participants jointly stress-test the taxonomy itself. By observing AI behaviours in a 'safe-to-fail' context, stakeholders can iteratively refine risk classifications and governance assumptions, ensuring that these remain evidence-based and technically grounded. In this way, sandboxes can help translate uncertainty into structured supervisory learning and narrow the gap between technological change and regulatory response. Functioning as risk observatories, they channel insights from experimentation back into regulatory guidance, supervisory practices, and, where appropriate, formal rule-making. Financial AI sandboxes may therefore support a transitional phase in which proportionate, evidence-based regulatory models are not only conceptualised, but progressively tested and ultimately institutionalised.[78]

[76] Cf. *Ghetti et al.* [36], who similarly argue that a MiFID III reform may be required to codify AI-specific, proportionate governance obligations.
[77] For an overview of regulatory sandboxes across selected global jurisdictions, see *McCarthy* [39]. For recent empirical data on sandbox initiatives within the EU, see *AFME* [1], p. 25.
[78] See ESAs [27].

4. Conclusions

MiFID II's commitment to technological neutrality once ensured adaptability, competition, and innovation. By regulating functions rather than tools, it allowed EU financial law to accommodate successive waves of digital transformation. Yet, as AI becomes embedded across all layers of financial activity, neutrality without differentiation risks producing normative silence. The Directive still captures AI within its scope but not always its distinctive risks—those arising from systems that learn, adapt, and act with limited human involvement.

The path forward is evolutionary rather than revolutionary. The principle of proportionality offers the bridge through which MiFID II can develop into a risk-based framework that aligns regulatory intensity with technological and systemic risk. A differentiated taxonomy—horizontal across domains and refined within them—could reconcile flexibility with accountability, preserving MiFID II's technology-agnostic spirit while embedding proportionate obligations that reflect the operational realities of learning systems. However, this evolution could unfold in a fragmented regulatory landscape. Supervisory capacity and enforcement cultures vary across Member States, while disparities in AI infrastructure and expertise may amplify divergence. Without coordination, these asymmetries risk deterring responsible adoption and concentrating innovation within jurisdictions with stronger supervisory and technological ecosystems.

The implications reach beyond mere regulatory compliance, as effective AI governance will be integral to the future credibility of the Capital Markets Union. A coherent, risk-based framework could provide the harmonised supervisory standards needed to ensure cross-border consistency and trust. Conversely, fragmented approaches may entrench regulatory silos and hinder integration, while also exposing EU capital markets to new forms of risks. Ultimately, governing financial AI is less about replacing old rules than reinterpreting them for evolving socio-technical realities—ensuring that Europe's financial system remains innovative, safe, resilient, and normatively robust in the age of non-human 'intelligent' systems.

5. References

AFME: Capital Markets Union. Key Performance Indicators – Seventh Edition: Unlocking Capital Markets for a Competitive Europe, November 2024 (2024). Available at: https://www.afme.eu/media/edmbujal/afmecmukpis2024061.pdf.

Aikman, D, Galesic, M., Gigerenzer, G., Kapadia, S., Katsikopoulos, K., Kothiyal, A., Murphy, E., Neumann, T.: Taking uncertainty seriously: simplicity versus complexity in financial regulation. Industrial and Corporate Change **30(2)**, pp. 317–345 (2021).

Annunziata, F.: Artificial Intelligence and Market Abuse Legislation: A European Perspective. Edward Elgar Publishing, Cheltenham (2023).

Appermont, N.: 'A conceptual framework on legal complexity'. The Theory and Practice of Legislation **13(2)**: pp. 236-263.

Azzutti, A.: AI Governance and Algorithmic Trading: Some Regulatory Insights from the EU AI Act. Banking & Finance Law Review **41(1)**, pp. 133-168 (2024).

Azzutti, A.: The Algorithmic Future of EU Market Conduct Supervision: A Preliminary Check. In: Böffel, L., Schürger, J. (eds.) Digitalisation, Sustainability and the Banking and Capital Union. Palgrave Macmillan, Cham (2023).

Azzutti, A.: AI Trading and the Limits of EU Law Enforcement in Deterring Market Manipulation. Computer Law & Security Review **45**, Article 105690 (2022).

Azzutti, A.; Magalhães Batista, P.; Ringe, W.-G.: Good Administration in AI-Enhanced Banking Supervision: A Risk-Based Approach. Columbia Journal of European Law **29(3)**, pp. 434-497 (2024).

Azzutti, A., Ringe, W.-G., Stiehl, H.-S.: Machine Learning, Market Manipulation and Collusion on Capital Markets: Why the "Black Box" Matters. University of Pennsylvania Journal of International Law **43**, pp. 79-135 (2021).

Barasa, M., di Castri, S., Grasser, M., Kiuhan, S., Letsiou, K., Sousa Faira, L.: State of SupTech Report 2024. Available at: https://ssrn.com/abstract=5518142.

BIS: Regulating AI in the financial sector: recent developments and main challenges. FSI Insights, No 63, 12 December 2024. Available at: https://www.bis.org/fsi/publ/insights63.pdf.

Black, J; Baldwin, R.: Really responsive risk-based regulation. Law and Policy **32(2)**, pp. 181-213.

Buczynski, W., Steffek, F., Jamnik, M., Cuzzolin, F., Sahakian, B.: Future themes in regulating artificial intelligence in investment management. Computer Law & Security Review **56**, Article 106111 (2025).

Buttigieg, C.P., Armeni Cauchi, A.: Going beyond the de Larosière Doctrine: effectiveness with proportionality. ERA Forum **26**, pp. 21-40 (2025).

Clausmeier, D.: Regulation of the European Parliament and the Council on digital operational resilience for the financial sector (DORA). International Cybersecurity Law Review **4**, pp. 79-90 (2023).

Della Corte, L: On proportionality in the data protection jurisprudence of the CJEU. International Data Privacy Law **12(4)**, pp. 259-275 (2022).

EASA: EASA Concept Paper: Guidance for Level 1 & 2 Machine Learning Applications. A Deliverable for the EASA AI Roadmap. March 2024. Issue 02. Available at: https://www.easa.europa.eu/en/downloads/139504/en.

EBA: Rising application of AI in EU banking and payments sector (2025). Available at: https://www.eba.europa.eu/sites/default/files/2025-09/146b3558-d026-47bf-a872-f05e93ed30d2/Rising%20application%20of%20AI%20in%20EU%20banking%20and%20payments%20sector.pdf.

EBA: Draft Regulatory Standards on Operational Risk Losses Mandates. EBA/RTS/2025/03, 04 August 2025. Available at: https://www.eba.europa.eu/sites/default/files/2025-08/1f9809f8-13bf-4365-aadc-4e7a8cdf89f1/Final%20Report%20on%20RTS%20Operational%20risk%20losses%20mandates.pdf.

EBA: Opinion of the European Banking Authority on money laundering and terrorist financing risks affecting the EU's financial sector. EBA/Op/2025/10, 28 July 2025. Available at: Opinion and Report on ML TF risks.pdf.

EBA: Special topic – Artificial intelligence (2024). Available at: https://www.eba.europa.eu/publications-and-media/publications/special-topic-artificial-intelligence.

EBA: Machine learning for IRB models. Follow-up report from the consultation on the discussion paper on machine learning for IRB models. EBA/REP/2023/28, August 2023. Available at: https://eba.europa.eu/sites/default/files/document_library/Publications/Reports/2023/1061483/Follow-up%20report%20on%20machine%20learning%20for%20IRB%20models.pdf.

EIOPA: Opinion on AI Governance and Risk Management. EIOPA-BoS-25-360 (06 August 2025). Available at: https://www.eiopa.europa.eu/document/download/88342342-a17f-4f88-842f-bf62c93012d6_en.

EMA: Reflection paper on the use of Artificial Intelligence (AI) in the medical product lifecycle (9 September 2024). Available at: https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf.

Epstein, R.A.: Simple Rules for a Complex World. Harvard University Press, Cambridge (1995).

ESAs: Joint Committee Update on Risks and Vulnerabilities in the EU Financial System – Spring 2025. Available at: https://www.esma.europa.eu/sites/default/files/2025-03/Joint_Committee_Update_on_risks_and_vulnerabilities_in_the_EU_financial_system_-_Spring_2025.pdf.

ESAs: Update on the functioning of innovation facilitators – innovation hubs and regulatory sandboxes. ESA 2023 27, 11 December 2023. Available at: https://www.esma.europa.eu/sites/default/files/2023-12/ESA_2023_27_Joint_ESAs_Report_on_Innovation_Facilitators_2023.pdf.

ESMA: Artificial Intelligence in EU Investment Funds: Adoption, Strategies and Portfolio Exposures. ESMA TRV Risk Analysis (Financial Innovation), 25 February 2025. Available at: https://www.esma.europa.eu/sites/default/files/2025-02/ESMA50-43599798-9923_TRV_Article_Artificial_intelligence_in_EU_investment_funds.pdf.

ESMA: Public Statement on the Use of Artificial Intelligence (AI) in the Provision of Retail Investment Services (30 May 2024). Available at: https://www.esma.europa.eu/sites/default/files/2024-05/ESMA35-335435667-5924__Public_Statement_on_AI_and_investment_services.pdf.

ESMA: Supervisory Briefing on Understanding the Definition of Advice under MiFID II (11 July 2023). Available at: https://www.esma.europa.eu/sites/default/files/2023-07/ESMA35-43-3861_Supervisory_briefing_on_understanding_the_definition_of_advice_under_MiFID_II.pdf.

ESMA: Guidelines on Certain Aspects of the MiFID II Suitability Requirements. ESMA35-43-3172, 03 April 2023. Available at: https://www.esma.europa.eu/sites/default/files/2023-04/ESMA35-43-3172_Guidelines_on_certain_aspects_of_the_MiFID_II_suitability_requirements.pdf.

ESMA: Artificial Intelligence in EU Securities Markets. ESMATRV Risk Analysis, 1 February 2023. Available at: https://www.esma.europa.eu/sites/default/files/library/ESMA50-164-6247-AI_in_securities_markets.pdf.

European Commission: 'Better regulation' toolbox 2023 (2023). Available at: https://commission.europa.eu/law/law-making-process/better-regulation/better-regulation-guidelines-and-toolbox/better-regulation-toolbox_en.

FSB: Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector (10 October 2025). Available at: https://www.fsb.org/uploads/P101025.pdf.

Gehrmann, S., Huang, C., Teng, X., Yuovski, S., Bhorkar, A., Thomas, N., Doucette, N., Rosenberg, D., Dredze, M., Rabinowitz, D.: Understanding and Mitigating Risks of Generative AI in Financial Services. In: FAccT '25: Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency. Association for Computing Machinery, New York (2025).

Ghetti, R., Novelli, C., Hacker, P., Floridi, L.: AI and Investment Services in EU Law: The Case for MiFID III. Available at: https://ssrn.com/abstract=5522839.

Kaplow, L.: Rules Versus Standards: An Economic Analysis. Duke Law Journal **42(3)**, pp. 557–629 (1992).

Leitner, G., Singh, J., van der Kraaij, A., Zsámboki, B.: The rise of artificial intelligence: benefits and risks for financial stability. In: ECB, Financial Stability Review, May 2024 (2024). Available at: https://www.ecb.europa.eu/press/financial-stability-publications/fsr/special/html/ecb.fsrart202405_02~58c3ce5246.en.html.

McCarthy, J.: From childish things: the evolving sandbox approach in the EU's regulation of financial technology. Law, Innovation and Technology **15(1)**: pp. 1-24.

MDCG: Guidance on Qualification and Classification of Software in Regulation (EU) 2017/745 – MDR and Regulation (EU) 2017/746 – IVDR (June 2025). Available at: https://health.ec.europa.eu/document/download/b45335c5-1679-4c71-a91c-fc7a4d37f12b_en.

MAS: Consultation Paper on Guidelines on Artificial Intelligence Risk Management. Consultation Paper, P017-2025 (November 2025): https://www.mas.gov.sg/-/media/mas-media-library/publications/consultations/bd/2025/final_consultation_paper_on_guidelines_on_ai_risk_management_forrelease.pdf.

Naeem, R., Treacy, S.: EU securities authority launches AI survey. Linklaters, 04 June 2025. Available at: https://financialregulation.linklaters.com/post/102kda6/eu-securities-authority-launches-ai-survey.

OECD: Regulatory approaches to Artificial Intelligence in finance. OECD Artificial Intelligence Papers, No. 24. OECD Publishing, Paris (2024). Available at: https://doi.org/10.1787/f1498c02-en.

Ojanen, A.: Technology neutrality as a way to future-proof regulation: The case of the Artificial Intelligence Act. European Journal of Risk Regulation, pp. 1-16 (2025).

Raschner, P.: Supervisory Oversight of the Use of AI and ML by Financial Market Participants. In: Böffel, L., Schürger, J. (eds.) Digitalisation, Sustainability and the Banking and Capital Union. Palgrave Macmillan, Cham (2023).

The RegTech Association and Deloitte: The RegTech Association Industry Perspectives 2024 (2024). Available at: https://regtechglobal.org/Industry-Perspectives.

Wendt, D.W.: AI Strategy and Security: A Roadmap for Secure, Responsible, and Resilient AI Adoption. Apress, Berkeley (2025).

Wittin, J., Ossowska, A., Sawicka, K.: Is DORA an opportunity for more balanced distribution of third-party risk management assurance responsibilities between banks, regulators and technology providers? Information & Communications Technology Law, pp. 1-8 (2025).

6. About the Authors

Dr Alessio Azzutti is a Lecturer in Law and Technology (FinTech) at the University of Glasgow. His research lies at the intersection of law, finance, and emerging technology, with a particular focus on artificial intelligence (AI). His published work examines AI governance and financial regulation, especially the ethical and legal challenges raised by the development and use of AI in financial services, including in the areas of RegTech and SupTech. He also leads innovative teaching in this field, including the microcredential course *AI & RegTech in Financial Compliance*. Before joining Glasgow, he held research positions at the National University of Singapore and Universität Hamburg. He holds degrees in Business Administration (B.A.), Finance and Risk Management (M.Sc.), and Law and Economics (LL.M.), as well as a Doctorate in Law (*cum laude*) from Universität Hamburg.

Get in touch

FRIL@FinTechScotland.com

FinTech Scotland®

University of Glasgow

University of Strathclyde Glasgow