Crypto & TradFi

AI Special: The Global Regulatory Acceleration

Unravelling the regulations for investors

Frontier AI: where financial stability meets global governance

Previous editions of the Regulatory Brief have outlined the European regulatory framework in successive layers: the AI Act and its Omnibus (Seqlense Regulatory Brief Issue #7), its relationship with the GDPR (Issue #8), the post-quantum transition (Issue #9), technological sovereignty (Issue #10), the convergence of AI and cybersecurity (Issue #11), dual supervision of market abuse and AML/CFT (Issue #12), the end of the MiCA grandfathering clause (Issue #13), and heightened vigilance regarding market abuse (Issue #14). This fifteenth edition continues this line of thought by focusing on an exceptional institutional moment: the first week of July 2026 saw the convergence of five major publications that redefine the role of artificial intelligence in the financial system and in the international order.

Between 6 and 7 July 2026, five major institutional initiatives followed one another in quick succession. In Europe, the ESRB issued a warning on systemic cyber risks linked to cutting-edge AI models, immediately backed by the three European supervisory authorities (EBA, EIOPA and ESMA). On the same day, the ECB called on significant institutions to strengthen their cyber defences without delay, with an action plan due by 31 October 2026. In the United Kingdom, the Bank of England and the FCA published two major analyses on the impact of AI on financial stability and financial services. Finally, in Geneva, the United Nations launched the first Global Dialogue on AI Governance.

This convergence reflects an unprecedented institutional shift. Whereas AI had previously been treated as a regulatory project requiring structure (the AI Act, DORA, sector-specific guidelines), it is now classified by the European Systemic Risk Board (ESRB), the European macro-prudential regulator, as a source of systemic risk, with the cyber risk rating set at ‘severe’ (compared with ‘high’ in March 2026). This acceleration is no longer a hypothesis: it is a fact, from which the authorities are drawing immediate operational conclusions.

For regulated entities such as, amongst others, credit institutions, asset management firms, investment firms, CASPs, fintechs and market infrastructures, as well as for investors, these publications call for a unified interpretation: the supervision of AI is no longer a specialised subject confined to a few technical teams, but a matter of financial stability addressed by the highest European, UK and global authorities.

The main signal

AI is moving beyond the realm of sector-specific regulation into that of financial stability: the European macro-prudential regulator now classifies cyber risk associated with cutting-edge AI models as 'severe'.

The most significant development of the first week of July 2026 is not the existence of an AI cyber risk, which has been known since the publication of ENISA's *Threat Landscape Finance* in February 2025 and reiterated by the AMF, the ESAs and the AFM in June 2026 (see Issue #11). Rather, it is the shift towards a macro-prudential classification of the issue. The ESRB, in its warning of 25 June, describes frontier AI models as constituting a 'paradigm shift' for cybersecurity and a source of 'systemic risks to the EU financial system'. This reclassification shifts the issue from the realm of individual operational risk (addressed by DORA) to that of systemic risk (addressed by the macroprudential framework), and requires national and European regulators to incorporate it into their supervisory work.

This has two sets of operational implications. For supervised institutions, the ECB has set a specific deadline: a cyber action plan must be submitted to their Joint Supervisory Team by 31 October 2026. For the authorities, the ESRB calls on the European Union to "develop its capacity, expertise and strategic autonomy" in this area, directly echoing the technological sovereignty initiatives described in Issue #10 (*Technological Sovereignty Package* of 3 June 2026). AI supervision thus aligns with the European sovereignty agenda: the same causes (dependence on a small number of non-European suppliers) call for the same responses (strengthening European capabilities).

Focus 1 - Europe: ESRB, ESAs and ECB in coordinated action

On 7 July 2026, three coordinated European publications reached the same assessment regarding frontier AI models (FAIMs). The ESRB published a warning dated 25 June 2026, accompanied by an analytical report entitled 'Addressing Frontier AI Models with cyber capabilities from a financial stability perspective'. The ESRB's General Board had at that time classified systemic cyber risk as 'severe' (compared with 'high' in March 2026).

The assessment is clear. Frontier AI models, defined by the ESRB as advanced general AI models capable of significantly affecting offensive or defensive cyber operations, are capable, according to current assessments, of discovering vulnerabilities, generating functional exploits and autonomously executing full-scale cyberattacks at a speed, scale and level of precision far superior to those of previous models. This development is described as a paradigm shift and a critical inflexion point in AI capabilities. In the short and medium term, these offensive capabilities will dominate, with defensive applications catching up. Furthermore, the ESRB emphasises that the concentration of the main providers of state-of-the-art AI models outside the European Union exposes the region to strategic dependence and geopolitical risk.

On the same day, 7 July, the three European supervisory authorities, the EBA, EIOPA and ESMA, issued a joint statement welcoming and supporting the warning. They point out that the European framework with DORA and the AI Act provides a solid basis for addressing ICT and AI risks, but stress that the rapid evolution of models requires financial institutions to adapt their cyber capabilities, and competent authorities to incorporate these risks into their supervisory and oversight work. As the supervisory authorities for critical third-party ICT providers under DORA, the ESAs also state that they are engaging with these providers regarding the measures they are taking.

Finally, on the same day, Claudia Buch, Chair of the ECB's Supervisory Board, sent a letter to the CEOs of all significant institutions supervised under the Single Supervisory Mechanism. In it, she requested a cyber action plan, including immediate priorities and medium-term strategic directions, to be submitted to the Joint Supervisory Teams (JSTs) by 31 October 2026. The ECB has also extended the annual collection of the *IT Risk Questionnaire* from September 2026 to February 2027 to allow for this new deadline. A horizontal analysis of the submitted plans will be carried out. On the same day, the Bank of England issued a similar warning in its *Half-Year Financial Stability Report*, in line with the joint statements by the BOE, the FCA and the UK Treasury on frontier models and operational resilience.

Key points to watch

• The **deadline of 31 October 2026** for Significant Institutions to submit their cyber action plans, and the ECB's subsequent horizontal analysis of these plans.

• The **expected extension to all non-significant financial institutions** by the relevant national authorities, such as the ACPR in France, BaFin in Germany, and the AFM and DNB in the Netherlands.

• The **concrete measures taken by critical third-party ICT providers** designated under DORA, in response to discussions with the ESAs.

Focus 2 - The FCA's Mills Review: the UK's bet on agentic finance

On 6 July 2026, the FCA published the Mills Review, a landmark review led by Sheldon Mills, the authority's Executive Director, at the request of its Board, on the transformation of UK retail finance by AI by 2030 and beyond. This review is presented as the first work of its kind initiated by a regulator. It draws on more than 140 written contributions, a survey by Yonder Consulting of 5,026 UK financial consumers, and the FCA's own internal work since the launch of the *engagement paper* in January 2026.

The central finding is that retail finance is shifting from a 'human-led, episodic' model to an 'AI-enabled, continuous and delegated' model. The Mills Review proposes a five-level 'autonomy spectrum', ranging from human-operator (Level 1) to human-observer (Level 5). It identifies four systemic shifts: the transformation of firms' operations; the evolution of customer journeys; the reshaping of competitive dynamics and market power; and the amplification of both threats (fraud, cyber) and defences. A striking figure: according to the survey, 20 per cent of UK consumers, equivalent to 11 million adults, say they are prepared to use AI acting autonomously within predefined objectives (agentic AI). On the regulatory front, the FCA confirms that it will not create a new, standalone regime for AI: it will apply existing frameworks, the Consumer Duty, the Senior Managers Regime and the operational resilience framework, and will publish an *AI Good and Poor Practice publication* later in 2026.

The Review sets out seven priority recommendations for the FCA Board: to secure and adapt the regulatory scope; to strengthen coordination and oversight at the system level; to monitor the transition to autonomous models and adapt regulatory frameworks accordingly; to develop the FCA's AI Lab; to lay the foundations for 'agentic finance' (identity, authority, accountability, execution); adopt an AI-enabled agentic supervisory model; and develop a public financial service for AI-enabled financial capability. The FCA's Chair, Ashley Alder, hails the review as a response to the acceleration of AI, based on the principles-based and results-oriented approach already adopted by the UK regulator.

Key points to watch

• The publication of the **FCA's 'AI Good and Poor Practice',** expected later in 2026, will be the UK regulator's first operational framework on acceptable uses of AI.

• The **review of the regulatory scope** is recommended within three to six months, focusing on general-purpose large language models (LLMs) outside the FCA's remit and their impact on the boundary between advice and guidance.

• The **potential comparison between the UK approach (application of existing frameworks) and the continental approach** (sector-specific AI Act), particularly for pan-European players.

Focus 3: The UN launches the Global Dialogue on AI Governance in Geneva

On 6 and 7 July 2026, the UN held the first session of the *Global Dialogue on AI Governance* at Palexpo in Geneva, an intergovernmental forum established by General Assembly Resolution A/RES/79/325. This is the first universal forum in which all UN Member States, not just the technological powers, are brought together to set out guidelines for the governance of artificial intelligence. The Dialogue is co-chaired by Ambassadors Egriselda López (El Salvador) and Rein Tammsaar (Estonia). The inaugural session was opened by the United Nations Secretary-General, António Guterres, and the President of the General Assembly, Annalena Baerbock. A second session is already scheduled to take place in New York in May 2027.

The Secretary-General set out four priorities: common safety standards; red lines on human rights; capacity-building for developing countries (proposal for a Global Fund for AI); and an AI Environmental Transparency Initiative. He also called for an AI Child Safety Pledge, stressing that "no child should be a guinea pig for unregulated AI". Annalena Baerbock emphasised the fight against the malicious use of AI, noting that a commonly cited figure suggests that 99 per cent of deepfakes are of a sexual nature and that 96 per cent target women and girls. Guterres also described lethal autonomous weapons as morally repugnant.

The Dialogue is supported by a second mechanism: the *Independent International Scientific Panel on Artificial Intelligence*, comprising 40 experts and co-chaired by Yoshua Bengio (Canada) and

Maria Ressa (Philippines). This panel published its first preliminary report on 1 July 2026, serving as a common scientific basis for the Dialogue's discussions. Its conclusions are unambiguous: science "cannot guarantee" that cutting-edge AI models will not cause catastrophic harm, and AI "is outpacing the rate at which scientific understanding and governments' capacity to adapt are progressing". For the financial sector and investors, this work confirms that AI has entered a phase of global governance, which aligns with and amplifies the European and UK regulatory dynamics described in Focus 1 and 2.

Key points to watch

• The **second session of the Global Dialogue on AI Governance** in New York in May 2027, and the operational conclusions likely to emerge from it.

• The **upcoming reports from the Independent International Scientific Panel** and their influence on national and regional regulatory decisions.

• The **proposals for the Global Fund for AI and the AI Environmental Transparency Initiative** put forward by the UN Secretary-General, and their potential translation into concrete measures.

Focus 4: The implementation of the AI Act is drawing nearer: content transparency and French oversight

The international developments described in the previous focuses come just a few weeks before several major milestones under the European AI Act. On 2 August 2026, the transparency obligations set out in Articles 50(1), (3) and (4) of the AI Act will come into force: chatbots, emotion recognition systems, deepfakes and AI-generated content must be labelled in accordance with the prescribed procedures. On 2 December 2026, the obligation to label AI-generated content within the meaning of Article 50(2) will apply to pre-existing systems, and the so-called 'nudification' ban will come into force. On 2 December 2027, the rules applicable to high-risk systems listed in Annex III (credit scoring, life and health insurance pricing, amongst others) will apply in full.

In this context, the Code of Practice on Transparency of AI-Generated Content, published on 10 June 2026, provides operational guidance for providers and deployers of generative AI systems. For the financial sector, the stakes are high: investment advertisements, finfluence content, market videos, newsletters, automated analyses, customer service chatbots, recommendations and sponsored messages may all be produced or amplified by AI. Transparency regarding the origin of content is becoming a key issue for investor protection, linked to the MAR and MiCA frameworks and national regulations (the AMF's General Regulations, ESMA's guidance on marketing communications).

At the same time, on 30 June 2026, the AMF published its 2026 Market and Risk Map (see Issue #14). It confirms that geopolitical and cyber risks remain at the forefront and explicitly establishes the link with the rapidly evolving threats associated with AI, referring to its alert of 3 June 2026 (see Issue #11). In the report, the AMF announces work on how listed companies are integrating AI, cybersecurity and geopolitics into their governance, a strong signal to equity investors, ESG investors and proxy advisers. These elements tie in with the ESAs' doctrine presented in Focus 1 and confirm that AI governance is gradually becoming a market issue in its own right.

Key points to watch

• Effective compliance by **2 August 2026** with the Article 50 transparency obligations by financial firms deploying chatbots, robo-advisers or AI-generated content intended for the public.

• Upcoming **AMF publications on AI governance in listed companies**, and their implications for ESG scoring frameworks and proxy voting.

• The **practical implementation of the Code of Practice on Transparency** by the main providers of generative models and financial media organisations.

Key takeaways

Five fundamental transformations are underway:

• AI is crossing a threshold: the **European Systemic Risk Board (ESRB) has classified the cyber risk associated with cutting-edge AI models as 'severe'**, bringing AI into the realm of financial stability, which is addressed at the level of the ESRB's General Board.

• The **European supervisory agenda is taking shape** with the 31 October 2026 deadline set by the ECB for the cyber action plans of *Significant Institutions*, and the launch of a phase of enhanced *oversight* of critical ICT providers by the ESAs.

• The **United Kingdom is embracing agentic finance** with the Mills Review of 6 July 2026: application of existing frameworks (Consumer Duty, Senior Managers Regime), a five-level autonomy spectrum, and seven recommendations including a model for AI-enabled agentic supervision.

• **Global AI governance is taking shape at the UN** with the Global Dialogue on AI Governance in Geneva (6-7 July 2026), the preliminary report from the Independent International Scientific Panel (1 July), and a second session scheduled for New York in May 2027.

• The **European AI Act enters its practical implementation phase** with the transparency obligations under Article 50 coming into force on 2 August 2026, labelling on 2 December 2026, and the entry into force of the high-risk rules (credit, insurance) on 2 December 2027, milestones now underpinned by an operational code of conduct published on 10 June.

Conclusion

The first week of July 2026 will be remembered as a rare moment of institutional convergence, when four centres of authority, Frankfurt (ESRB, ECB), Brussels (ESAs), London (FCA, Bank of England) and Geneva (UN), simultaneously conveyed the same message: AI is no longer a regulatory project to be structured, but a systemic reality to be supervised. For regulated entities, the operational implication is clear. Cyber security measures must be reviewed in light of the risk level now classified as 'severe' by the ESRB, within a tight timeframe (31 October 2026 for payment systems supervised by the ECB). AI governance must be elevated to the level of the Executive Committee and the Board, in line with the approach set out in the Mills Review and by the AMF. The transparency requirements of the AI Act must be operational by 2 August 2026. For investors, issuers' AI governance is becoming a criterion for analysis, to be incorporated into ESG frameworks and proxy voting exercises. And for all stakeholders, the global dimension of the issue, launched by the Global Dialogue on AI Governance, serves as a reminder that national and regional frameworks are now part of an international architecture currently under development.

More broadly, this development confirms that the governance of artificial intelligence is gradually becoming a new pillar of international financial regulation. For both financial institutions and investors, the ability to anticipate these developments will now constitute a genuine strategic advantage.

Main sources

- **ESRB (European Systemic Risk Board)**, *Warning on Systemic Cyber Risks Stemming from Frontier Artificial Intelligence Models*, warning dated 25 June 2026 and published on 7 July 2026; accompanying report, *Addressing Frontier AI Models with Cyber Capabilities from a Financial Stability Perspective* (esrb.europa.eu).
- **ESAs (EBA, EIOPA, ESMA)**, *The ESAs Support ESRB Warning on Systemic Cyber Risks from Frontier AI Models*, joint press release, 7 July 2026 (eba.europa.eu, eiopa.europa.eu, esma.europa.eu).
- **European Central Bank (ECB)**, Claudia Buch (Chair of the Supervisory Board), letter to the CEOs of significant institutions supervised under the Single Supervisory Mechanism (SSM), 7 July 2026; extension of the 2026 IT Risk Questionnaire from September 2026 to February 2027 (bankingsupervision.europa.eu).
- **Bank of England**, *Financial Stability Report - July 2026*, section on frontier AI models and operational resilience.
- **Financial Conduct Authority (FCA)**, *The Mills Review - AI and the Future of Retail Financial Services*, published on 6 July 2026 under the direction of Sheldon Mills (Executive Director, Consumers and Competition), commissioned by the FCA Board chaired by Ashley Alder (fca.org.uk).
- **Financial Conduct Authority (FCA)**, *Yonder Consulting AI Survey*, survey of 5,026 UK adults conducted in April 2026 and published as an appendix to *The Mills Review*.
- **United Nations**, *Global Dialogue on AI Governance - First Session*, Geneva (Palexpo), 6-7 July 2026; established pursuant to UN General Assembly Resolution A/RES/79/325; opening addresses by António Guterres and Annalena Baerbock (un.org).
- **United Nations**, *Independent International Scientific Panel on Artificial Intelligence - Preliminary Report*, 1 July 2026, co-chaired by Yoshua Bengio and Maria Ressa.
- **European Commission**, *Code of Practice on Transparency of AI-Generated Content*, published on 10 June 2026 (digital-strategy.ec.europa.eu).
- **Autorité des marchés financiers (AMF)**, *2026 Market and Risk Outlook* (*Cartographie 2026 des marchés et des risques*), published on 30 June 2026 (amf-france.org).
- **Regulation (EU) 2024/1689 (AI Act)**, in particular Articles 50(1), 50(2), 50(3), 50(4), and Annex III; application milestones: 2 August 2026, 2 December 2026, and 2 December 2027.
- **Regulation (EU) 2022/2554 (DORA)**, applicable since 17 January 2025; European framework for digital operational resilience.
- **Additional industry analyses consulted:** A&O Shearman, CMS, Bratby Law, TLT, Regulation Tomorrow, Insight EU Monitoring, Finadium, PYMNTS, Retail Banker International, ThePlanetTools.ai, and UN News (UN Web TV/Unifeed)

.***