

Crypto, TradFi & Digital

Grey areas: ESMA maps hybrid products, the AI Act tests its incident notification

Decoding financial, AI and data regulation to anticipate, comply and decide

Prediction markets, tokenised equities, an AI incident: the European frameworks face their first borderline cases

Previous editions of the Regulatory Brief have documented the ramp-up of the major European frameworks: the full application of MiCA (Edition #13), the AI Act as at 2 August 2026 (Edition #18), the clarification of the perimeter of advice on crypto-assets (Edition #20), the review of the MiCA framework and the Pontes/Appia infrastructure (Edition #22), and the entry into the fine mechanics of enforcement (Edition #23). This twenty-fourth edition continues that trajectory by observing the European frameworks at the precise point where they are most severely tested: their boundaries.

Two publications this week illustrate this moment. On 10 September 2026, ESMA published its second TRV Risk Monitor of the year, which devotes two in-depth analyses to prediction markets and tokenised equities, two families of products whose legal characterisation determines the entire applicable regime (MiFID II, MiCA, MAR, or even national gambling law). On 7 September 2026, the European Commission confirmed that it had received from OpenAI a report on the unauthorised use of a German website by AI agents during experiments: one of the first public tests of the incident-notification mechanism provided for by the AI Act for general-purpose AI models presenting systemic risk.

These two cases, very different in their subject matter, raise the same question: that of characterisation. For prediction markets and tokenised equities, the technology used (blockchain, settlement stablecoin) does not on its own determine the applicable law; it is the characteristics of the contract and the rights attached to the token that determine the regime. For the incident reported by OpenAI, it is the legal characterisations still open (systemic-risk model or not, serious incident or experimental anomaly, binding obligation or cooperative step) that will determine the reach of the precedent. In both cases, the boundary of the framework is where its effectiveness is decided.

The main signal

The maturity of the European frameworks is now measured at their boundaries: the legal characterisation of a hybrid product or an AI incident becomes the decisive act from which everything else follows, namely the applicable regime, the competent authority, the obligations and the sanctions.

ESMA's report shows this precisely for hybrid financial products. A prediction-market contract may constitute a financial instrument within the meaning of MiFID II (generally a derivative), fall under MiCA where it relies on DLT without constituting a financial instrument, or be treated as a gambling product under national law. A “tokenised equity” may be a share genuinely issued on DLT (native issuance) or a token representing a mere economic exposure to a share held off-chain by an intermediary (a so-called “wrapped” structure). In each case, the characterisation commands the regime, and the commercial label says nothing about the applicable law.

The OpenAI case illustrates the same logic on the AI side. Article 55 of the AI Act applies only to providers of general-purpose AI models presenting systemic risk. The public information does not, at this stage, make it possible to know whether the model involved falls within that category, whether the report was submitted under a binding obligation or as a cooperative step, or whether the facts have been formally characterised as a serious incident. The way the AI Office handles this case will, de facto, set the first practical contours of these characterisations. For regulated players in both universes, the lesson is identical: the characterisation analysis, often treated as a technical preliminary, is becoming the core of the compliance work.

Focus 1: Prediction markets and tokenised equities, ESMA maps the new boundaries between MiFID II, MiCA and MAR

In its second TRV Risk Monitor of 2026, published on 10 September, the European Securities and Markets Authority devotes two analyses to prediction markets and tokenised equities. The report creates no new obligation, but identifies risks that may require heightened vigilance from authorities and market players.

Prediction markets: a legal characterisation still highly product-dependent

Prediction markets allow the trading of contracts whose payout depends on the occurrence of a future event: an election result, an economic indicator, a sporting event, the movement of an asset or a weather phenomenon. Activity has grown strongly since the end of 2024. In the fourth quarter of 2025, quarterly volumes stood at around 8.8 billion dollars on Kalshi and 12 billion on Polymarket. ESMA nevertheless points out that these figures are mainly global and do not make it possible to assess precisely the participation of European retail investors.

In the European Union, the characterisation depends on the specific features of each contract. It may constitute a financial instrument within the meaning of MiFID II, generally a derivative, where the event relates to one of the underlyings listed in Annex I of the directive. Where it relies on DLT without constituting a financial instrument, it may fall under MiCA. It may also be treated as a gambling or betting product under national law. Where a contract is characterised as a financial instrument, its marketing may also be subject to the national intervention measures applicable to binary options, which generally prohibit their sale to retail investors. ESMA notes that the main international platforms do not currently hold the European authorisation required to provide regulated services in the Union. It would therefore be reductive to present these markets as simply a new category of crypto-assets: the use of a blockchain or a stablecoin to trade and settle the contract does not, on its own, determine the applicable regulation.

Market abuse, pseudonymity and oracles

ESMA identifies particular risks of market manipulation, insider dealing and wash trading. On platforms using pseudonymous wallets, a single operator may control several accounts, which complicates the detection of coordinated transactions and the reconstruction of positions. The platforms' surveillance mechanisms may, moreover, intervene too late, after the event has occurred and the winnings have been paid out. The Market Abuse Regulation (MAR) can help sanction such

conduct, but only where the contract concerned actually falls within the perimeter of financial regulation.

The determination of the event triggering the payout is another sensitive point. An ambiguous definition, a failing or manipulated data source, or an opaque platform decision can lead to a disputed settlement. In decentralised arrangements, these risks are heightened by the dependence on oracles and by the difficulty of reversing operations executed by smart contracts. ESMA finally stresses that AI-generated signals, algorithmic strategies and bots can accentuate information asymmetries, amplify volatility or facilitate coordinated behaviour. Misleading AI-generated content circulated on social media can also artificially alter the perceived probability of an event.

Tokenised equities: economic exposure does not mean ownership of the share

The tokenised-equity market remains very small, but is growing rapidly. According to the data cited by ESMA, the value of assets in circulation is thought to have risen from around 300 million euros at the end of 2024 to 1.9 billion euros at the end of June 2026, a 6.5-fold increase in eighteen months. Activity remains concentrated on a few large US technology companies and is still negligible compared with the capitalisation of traditional equity markets.

Above all, ESMA distinguishes two models: “native” issuance, in which the share itself is issued and recorded on DLT; and so-called “wrapped” structures, in which the token represents a claim on, or an economic exposure to, a share held off-chain by an intermediary. The second model currently predominates. Yet in these structures, the transfer of the token does not necessarily transfer legal title to the underlying share. The investor's rights then depend on the contractual documentation and the structure chosen. Particular attention should be paid to the existence of voting or financial rights, the arrangements for paying dividends and the consequences of a failure of the platform or the custodian. The interposition of a tokenisation platform and a custodian may therefore add new counterparty dependencies instead of removing intermediaries.

Fragmented liquidity and still-hybrid settlement

The expected benefits of tokenisation (programmability, automation of corporate actions, extended hours or atomic settlement) are still only partly realised. In many cases, the token is transferred on the blockchain while the payment is still made through conventional infrastructure. This hybrid architecture requires reconciliation between the two systems and limits the reduction of settlement risk. The coexistence of several non-fungible representations of the same share can also fragment liquidity across the different platforms. ESMA sees this as a risk for price formation, market depth and execution quality, particularly in periods of stress. Finally, an error in the code of a smart

contract can lead to an erroneous transfer or a misallocation of rights, which is difficult to correct because of the immutability of the blockchain.

The issue: characterising the right represented, not just the technology

An asset commercially presented as a “tokenised equity” does not automatically fall under MiCA. MiCA excludes from its scope crypto-assets characterised as financial instruments. The characterisation must therefore be based on the rights attached to the token and on its economic substance, in accordance with MiFID II and ESMA's guidelines. For the players concerned, the main control points are the legal characterisation of the product, the identity and soundness of the intermediaries, the exact nature of the investor's rights, the custody and settlement arrangements, and the market-abuse surveillance mechanisms. One reading precaution is also required: the report is a risk analysis calling for continuous monitoring; it does not, at this stage, propose a new European regime for prediction markets or tokenised equities.

What to watch

- The **follow-up given by national authorities** to ESMA's findings, in particular regarding the absence of European authorisation of the main international prediction-market platforms.
- The **interplay with the ongoing MiCA consultation** (see Edition #22), which questions precisely the treatment of prediction markets, perpetual contracts and the boundary with tokenised financial instruments.
- The development of the **tokenised-equity market** (volumes, native structures versus wrapped structures, the possible entry of European issuers) and the first supervisory clarifications on holders' rights.

Focus 2: AI Act, OpenAI's report tests the European incident-notification regime

On 7 September, the European Commission confirmed that it had received from OpenAI a report on the unauthorised use of a German website by AI agents. According to the Reuters investigation, these agents are said to have used the site as a bulletin board to exchange information during experiments conducted in the spring. The Commission's spokesperson, Thomas Regnier, stressed that an incident report should not be a mere formality and should set out precisely the facts and the measures envisaged. He did not, however, indicate on what date OpenAI had informed the Commission; the Commission merely stated that it remained in close contact with the company. OpenAI has, for its part, publicly acknowledged the existence of this “wiki incident” and considered it necessary to define more clearly the arrangements for disclosing unexpected or misaligned behaviour observed during experiments.

The Article 55 regime does not concern all AI systems

Article 55 of the AI Act applies specifically to providers of general-purpose AI models presenting systemic risk. It requires them in particular to assess and mitigate systemic risks, carry out appropriate model evaluations, track, document and report serious incidents, communicate any corrective measures, and ensure an appropriate level of cybersecurity for the model and its infrastructure. These obligations have applied since 2 August 2025. By contrast, the Commission's supervisory and sanctioning powers over providers of general-purpose AI models have applied since 2 August 2026 (see Edition #18).

In November 2025, the Commission published a reporting template for serious incidents involving general-purpose AI models presenting systemic risk. These documents must be submitted to the AI Office and, where applicable, to the competent national authorities via the secure EU SEND platform. OpenAI is also among the signatories of the Code of Practice for general-purpose AI models. This code is a voluntary instrument for demonstrating compliance with Articles 53 and 55; signing it does not, however, mean that every OpenAI model is legally characterised as a model presenting systemic risk.

What the public information does not allow us to assert

Several elements remain undetermined at this stage. The Commission has not specified which OpenAI model was involved; whether that model was legally regarded as presenting systemic risk; whether the report was submitted under a binding Article 55 obligation or as part of a cooperative

step; whether the facts had been formally characterised as a serious incident; or whether an investigation or sanction procedure had been opened. Nor does the experimental context automatically settle the applicability of the regulation: while Article 2 of the AI Act provides a general exclusion for certain research, testing or development activities prior to placing on the market, the Commission points out that certain obligations relating to systemic-risk models may already concern their development phase where they are intended to be placed on the market.

A concrete test for the supervision of advanced models

The case raises several questions to which the AI Office's practice will progressively have to provide answers: how to distinguish an experimental anomaly from a serious incident that must be reported; at what point the provider has sufficient information to notify the authority; how far it must document the technical causes, the external effects and the autonomous behaviour observed; what corrective measures are expected where the incident results from a combination of the model, the agentic architecture and the test environment; and how to articulate the regulatory notification with the publication of information for the persons or organisations affected.

Since 2 August 2026, the AI Office has been able to request information and documentation, carry out evaluations, require corrective measures and sanction breaches of the obligations applicable to providers of general-purpose AI models. Nothing indicates, however, at this stage that these powers have been formally deployed against OpenAI in this case. A reading precaution is therefore required: one should speak of a report submitted to the Commission and of ongoing follow-up. The public elements do not make it possible to present the event as a serious incident formally characterised under Article 55, or as the opening of infringement proceedings.

What to watch

- The **follow-up given by the AI Office**: requests for additional information, evaluation of the model concerned, possible corrective measures required.
- The **emerging doctrine on the boundary between experimental anomaly and serious incident**, and on the applicability of the obligations to the development phases of models intended for the market.
- The **development of model providers' disclosure practices**: OpenAI has publicly considered it necessary to define more clearly the arrangements for disclosing unexpected behaviour observed during experiments, a sign that the transparency standards of the laboratories are taking shape.

Key takeaways

Five structural points to take on board:

- The **ESMA TRV Risk Monitor of 10 September 2026** maps the risks of prediction markets and tokenised equities without creating any new obligation: the characterisation of each product (MiFID II, MiCA or national gambling law) remains the decisive act.
- The **prediction markets are growing strongly** (around 8.8 billion dollars of quarterly volumes on Kalshi and 12 billion on Polymarket in Q4 2025, global figures) but the main international platforms do not hold the European authorisation needed to provide regulated services in the Union.
- The market in **tokenised equities has grown 6.5-fold in eighteen months** (from 300 million euros at the end of 2024 to 1.9 billion at the end of June 2026) but remains dominated by “wrapped” structures in which the transfer of the token does not necessarily transfer legal ownership of the share: the investor's real rights depend on the contractual documentation.
- The European Commission confirmed on 7 September 2026 that it had received from **OpenAI an incident report** on the unauthorised use of a German website by AI agents: one of the first public tests of the notification mechanism provided for by the AI Act, whose handling by the AI Office will help set the first practical standards.
- The **legal characterisations remain open in the OpenAI case**: neither the systemic character of the model involved, nor the formal characterisation as a serious incident, nor the binding or voluntary nature of the submission is publicly established. No sanction procedure is documented at this stage.

Conclusion

The two cases in this edition confirm a development the Regulatory Brief has been documenting for several months: the European frameworks are now judged at their boundaries. For the hybrid products analysed by ESMA, the boundary is that of characterisation: a prediction-market contract or a tokenised equity exists legally only through the regime it falls under, and the commercial label or the technology used prejudices nothing. For the AI Act, the boundary is that of notification: the distinction between experimental anomaly and serious incident, between ordinary model and systemic-risk model, between voluntary cooperation and binding obligation, will be built through the AI Office's practice, case by case. For regulated players, the operational conclusion is the same in both universes: invest in the characterisation analysis upstream, document the choices made, and follow closely the first decisions that will turn these open questions into established doctrine.

Main sources

- ESMA, *TRV Risk Monitor No. 2, 2026*, published 10 September 2026, in particular the studies “Tokenisation of equities” and “Prediction markets” (esma.europa.eu).
- Regulation (EU) 2023/1114 (MiCA); Directive 2014/65/EU (MiFID II), in particular Annex I; Regulation (EU) No 596/2014 (MAR).
- Reuters, “*OpenAI has sent EU incident report on hijacked German website, Commission says*”, 7 September 2026.
- European Commission, *General-Purpose AI Models in the AI Act, Questions & Answers* (digital-strategy.ec.europa.eu).
- European Commission, *Reporting template for serious incidents involving general-purpose AI models with systemic risk*, 4 November 2025; secure EU SEND platform.
- Regulation (EU) 2024/1689 (AI Act), in particular Articles 2, 51, 52 and 55; Code of Practice for general-purpose AI models.

The Seqlense Regulatory Brief - Crypto, TradFi & Digital - Edition #24

This publication is provided for information purposes only and constitutes neither legal advice, nor investment advice, nor a personal recommendation. It cannot replace a legal or operational analysis specific to each organisation, taking account of its size, sector, practices and contractual ecosystem.

The information presented reflects a general analysis of regulatory dynamics as at the date of publication. The positions of European, American, Asian and national authorities may change rapidly; texts still under construction do not create any immediate obligation and must not be treated as applicable law.

Despite the care taken in selecting and verifying sources, no guarantee is given as to the accuracy, completeness or currency of the information. Operational practices should be adjusted in the light of forthcoming supervisory positions, enforcement decisions and the development of the legal frameworks.

Any compliance decision is the sole responsibility of the organisation concerned and should, where appropriate, be taken with the support of qualified legal and technical advisers in each relevant jurisdiction.