Quaderni AISDUE - Rivista quadrimestrale
ISSN 2975-2698 - Fascicolo n. 1/2026

BlogDUE

Taming systemic risk: The EU's GPAI Code of Practice between co-regulation & fundamental rights

Rosa Maria Torraco (MA Candidate at the European University Institute) – 31 gennaio 2026

SUMMARY: 1. Introduction. – 2. The AIA definition of GPAI models. – 3. systemic risk and the protection of fundamental rights in the AIA framework. – 4. From the AIA to the Code of Practice. – 5. Between innovation and the protection of fundamental rights. – 6. Final reflections and conclusions.

1. The Artificial Intelligence Act (Regulation (EU) 2024/1689, hereinafter "AIA") has quickly become a cornerstone of the European Union (EU)'s digital rulebook (there are numerous scholarly contributions on the AIA which address specific aspects of the Regulation). Among those examining the Regulation as a whole, without any claim to exhaustiveness, the following may be mentioned: L. COTINO HUESO, D. U. GALETTA, *The European Union Artificial Intelligence Act: A Systematic Commentary*, Napoli, 2025; A. HUERGO LORA, G. M. DÍAZ GONZÁLEZ (eds.), *The EU Regulation on Artificial Intelligence: a Commentary*, Milano, 2025; C. N. PEHLIVAN, N. FORGÓ, P. VALCKE (eds.), *The EU Artificial Intelligence (AI) Act: A Commentary*, Alphen aan den Rijn, 2024; P. VOIGT, N. HULLEN, *The EU AI Act Answers to Frequently Asked Questions*, Berlin, 2024). Its ambition is clear: to bring coherence to the development and use of AI across the Union, while keeping innovation anchored in the EU Charter of Fundamental Rights (the "Charter") and the founding values of the EU.

At the same time, the AI landscape is constantly evolving, and when the Commission first proposed the AIA in 2021, the idea of "general-purpose AI" (GPAI) was not yet on agenda. Only the Council's compromise text of 29 November 2021 included for the first time the category of "general purpose" AI systems (N. T. NIKOLINAKOS, *EU Policy and Legal Framework for Artificial Intelligence, Robotics and Related Technologies- The AI Act*, Cham, 2023). Following the launch of ChatGPT, the rapid rise of large language models and the intense public discussion around them prompted the European

Parliament to incorporate specific rules on general-purpose AI into its negotiating position, approved in June 2023 (G. OLIVATO, *Paving the Path Towards General Purpose AI Systems Regulation in the AI Act: an Analysis of the Parliament's and Council's Proposals*, in *MediaLaws*, 2024, p. 50 ff.). A final compromise was reached during the trilogue on 9 December 2023, where the Council, Commission, and Parliament agreed, *inter alia*, on the provisions regarding GPAI (E. PANDOLFI, *AI Act e pacing problem: la disciplina dell'IA per finalità generali,* in this *Journal*, 2025, p. 8).

To operationalize these provisions, the AIA called for the creation of codes of practice. The process came to a close on July 10, 2025, with the publication of the General-Purpose AI Code of Practice (the "Code"), which took effect less than a month later, on August 2, 2025, a date that also marks the entry into force of many provisions of the AIA. In parallel, given the interpretative challenges raised by the AIA provisions concerning GPAI models, the Commission adopted its Guidelines on the scope of the obligations for general-purpose AI models established by Regulation (EU) 2024/1689 (AI Act) (the "Guidelines"). Though voluntary, the Code sets out a demanding framework. Building on the hybrid nature of the AIA, it is intended not only to support compliance with the AIA s safety standards but also to strengthen the internal market, safeguard fundamental rights and democracy, and create space for innovation that remains trustworthy and human centred.

In this setting, the recent Digital Omnibus on AI Regulation Proposal (COM (2025) 836) does not alter the substantive obligations applicable to GPAI models in the AIA nor the scope of application of the Code, limiting itself to reinforce the supervisory and enforcement role of the AI Office in respect to GPAI systems.

The objective of this paper is to analyse the main features of the Code and place them within the broader framework of the AIA's obligations on GPAI models and systemic risk, while also considering their role in the wider context of EU law and its underlying values and objectives. This paper proceeds as follows. Section 2 traces the emergence of GPAI as a distinct regulatory category in the AIA. Section 3 analyses the concept of systemic risk within the AIA framework, situating it in relation to the Digital Services Act (Regulation 2022(EU) 2065, hereinafter "DSA") – another core Regulation in the EU digital sphere which contains well-established obligations on the mitigation and assessment of systemic risk – and showing how it connects market regulation with the protection of fundamental rights. Section 4 offers an overview of the AIA obligations applicable to GPAI models and of the Code, with its main pillars on transparency, copyright, and safety and security. Section 5 critically assesses the Code's co-regulatory nature and its attempt to balance innovation and competitiveness in the EU with fundamental rights protection. Finally, section 6 wraps up the conversation drawing some final reflections on this paper.

2. The AIA provides a legal definition of GPAI model in Art. 3(63), describing it as "an AI model, including where such an AI model is trained

with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market". This definition did not arise in isolation. It reflects the challenge of adapting the AIA's original risk-based framework, which was grounded in the notion of "*intended purpose*", to models whose versatility prevents any prior determination of specific uses. More broadly, it exposes the limits of a centralized, risk-based taxonomy when confronted with technologies whose risk profile is not only intrinsic to the system, but also generated downstream, by deployers and end users rather than by the model as such. Therefore, the idea of the legislator was to create an autonomous category of GPAI models, situated upstream of the classification of AI systems (E. PANDOLFI, *op. cit.*, p. 8).

Nevertheless, some scholars have expressed concerns about this definition from multiple angles (cfr., I. TRIGUERO ET AL., *General Purpose Artificial Intelligence Systems (GPAIS): Properties, Definition, Taxonomy, Societal Implications and Responsible Governance*, in *Information Fusion*, 2024, p. 2 ff.). Among these, the reference to "distinct tasks" in particular has been criticised as insufficiently precise and likely to require further legal interpretation (R. UUK, C. I. GUTIERREZ, A. TAMKIN, *Operationalising the Definition of General-Purpose AI: Assessing Four Approaches*, in *Law, Innovation and Technology*, 2025, p. 86 ff.). At the same time, it has been noted that the openness of the wording appears to have provided the necessary flexibility for the EU legislators to reach a political agreement (O. J. GSTREIN, N. HALEEM, A. ZEWITTER, *General-Purpose AI Regulation and the European Union AI Act*, in *Internet Policy Review*, 2024, p. 1 ff.).

The definition of GPAI models is not merely a theoretical matter but is bound to have practical implications on the providers of these models. For instance, an excessively broad definition of GPAI models could place disproportionate burdens even on providers of systems that in practice present only limited risks. For this reason, the recently issued Guidelines have sought to introduce further clarifications. They acknowledge that it is impossible to provide an exhaustive list of tasks or capabilities that would qualify a model as general-purpose, drawing in particular on Recitals No. 98 and 99 of the AIA, which point to model scale and large generative AI models as key indicators. To operationalise the concept, the Guidelines introduce a quantitative threshold of 10^23FLOPs, while at the same time recognising that this benchmark cannot serve as an absolute criterion. Indeed, while numerical thresholds may represent useful proxy, they risk reducing the complex interplay between AI capabilities and societal impact to a single technical proxy. This would be problematic from a rights perspective: smaller models could still spread disinformation, amplify discrimination, or jeopardise privacy, falling outside the perimeter of the corresponding AIA obligations.

Having provided a definition of GPAI models, the analysis will now turn to the notion of systemic risk, which, under the AIA, is intrinsically linked to such models.

3. The notion of systemic risk is not new to EU law, having long been recognized in contexts such as environment and cybersecurity (cfr., P. HACKER, A. KASIRZADEH, L. EDWARDS, *AI, Digital Platforms and the new Systemic Risk*, in *SSRN*, 2025; E. CIRONE, *L'AI Act e l'obiettivo (mancato?) di promuovere uno standard globale per la tutela dei diritti fondamentali*, in this *Journal*, No. 2, 2024, p. 51 ff.; V. ACHARYA, S. STEFFEN, *Analyzing Systemic Risk of the European Banking Sector*, in J.-P. FOUQUE, J. A. LANGSAM (eds.), *Handbook on Systemic Risk*, Barcelona, 2013, p. 247 ff.) and, more recently, in the DSA. It is precisely in relation to the latter that the concept of systemic risk in the AIA will be examined, with a view to clarifying the regulatory logic underpinning its adoption within the AIA (cfr., T. KARATHANASIS, *The Regulatory Interplay Between the AI Act and the DSA: Challenges, Burden, and Rationalization for AI Innovation in the EU*, in *SSRN*, 1st July 2025).

Unlike the AIA, the DSA does not provide a single definition of systemic risk (G. CONTISSA, F. GALLI, *La governance della "disinformazione aumentata" tra Digital Services Act e AI Act*, in *federalismi.it*, 30 May 2025, p. 126 ff.; G. ZACCARONI, *Intelligenza artificiale e principio democratico: riflessioni a margine dell'emersione di un quadro normativo europeo*, in this *Journal*, No. 2, 2024, p. 19 ff.). Instead, Art. 34(1) DSA divides the concept into four areas, addressing the spread of illegal content, potential infringements on fundamental rights such as freedom of expression and media pluralism, the erosion of civic discourse together with risks to electoral integrity and public security, and finally concerns related to gender-based violence, public health and well-being, and the protection of minors. In contrast, the AIA provides a specific legal definition in Art. 3(65), identifying systemic risk as a "risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the EU market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain". The AIA definition highlights three characteristics of systemic risk: it is linked to high-impact capabilities, it entails a significant impact on the EU market, such impact can be propagated at scale across the value chain. In this sense, the AIA specifies the nature and mechanics of systemic risk, but not its substantive types. By contrast, the DSA does exactly the opposite, providing a typology of risks without detailing their underlying nature or propagation mechanisms.

In other terms, while the protection of fundamental rights is explicitly prioritized, the AIA pursues this objective not through a platform-specific focus, as in the DSA, but through horizontal rules that extend across multiple sectors. This reflects a shift from regulating the online environment to

addressing the cross-cutting challenges posed by powerful AI systems whose negative effects may propagate at scale throughout the value chain. The two instruments should therefore be read as complementary rather than conflicting. The DSA addresses systemic risks within platform-based environments, whereas the AIA extends the concept to a horizontal level, targeting the wider challenges generated by general-purpose AI models. This relationship is confirmed by Recital No. 118 AIA, which stipulates that, where safeguards at model level prove insufficient to control systemic risks that the DSA cannot effectively address, the direct obligations of the AIA apply.

Accordingly, the AIA's treatment of systemic risk forms part of a broader continuum in EU regulatory practice. It cannot be confined to the sphere of product safety, but instead represents an attempt to combine market regulation with the protection of fundamental rights across interdependent digital domains (cfr., M. GORNET, W. MAXWELL, *The European Approach to Regulating AI Through Technical Standards*, in *Internet Policy Review*, 2024, p. 1 ff.; M. ALMADA, N. PETIT, *The EU AI Act: A Medley of Product Safety and Fundamental Rights?*, in *Robert Schuman Centre for Advanced Studies Research Paper*, 2023, p. 7 ff.). Furthermore, by adopting a rebuttable threshold in Art. 51(2), the AIA translates the principle of proportionality into a measurable governance mechanism. This ensures that only those GPAI models posing a real threat to rights and democracy fall within the systemic-risk regime. In this sense, technical thresholds operate not as neutral risk metrics, but as juridical filters that transform constitutional principles into enforceable obligations.

Art. 51 provides a clear illustration of this dynamic. Two distinct routes are envisaged for systemic-risk classification. The first one, set out in Art. 51(1)(a), applies where the model demonstrates what the Regulation describes as "high-impact capabilities", defined by Art. 3(64) AIA as "capabilities that match or exceed those recorded in the most advanced models". According to Art. 51(1)(a), this determination must be made on the basis of appropriate technical tools and methodologies. As specified by the Guidelines, such instruments will be further specified by delegated acts by the Commission. Nevertheless, a useful criterion is spelt out in Art. 51(2), which states that a GPAI model is "presumed to have high-impact capabilities […] when the cumulative amount of computation used for its training measured in floating point is greater than 10^25[FLOPS]". This computational threshold provides a concrete and measurable proxy for identifying models that are likely to have significant capabilities. At the same time, it functions as a rebuttable presumption rather than as an absolute rule, since models below the threshold may still fall within the definition of systemic risk where their capabilities or effects justify such a classification. However, critics argue that this criterion is problematic: as systemic risks depend on factors beyond compute, smaller models can be equally risky, and the threshold itself reflects political compromise designed to shield European firms such as Mistral and Aleph Alpha (M. EBERS, *Truly Risk-Based Regulation of Artificial Intelligence: How*

to Implement the EU's AI Act, in *Stanford–Vienna Transatlantic Technology Law Forum – European Union Law Working Papers*, No. 101, 2024, p. 31 ff.).

The second route, provided in Art. 51(1)(b), gives the Commission – on its own initiative or following a qualified alert from the scientific panel – the authority to classify a model as presenting systemic risk even where it does not fall neatly within the indicators specified under Art. 51(1)(a). This mechanism reflects the recognition that fixed technical benchmarks cannot exhaustively capture the evolving risks associated with GPAI, and that regulatory authorities require flexibility to address unforeseen developments.

This dual orientation confirms that the AIA cannot be reduced to a mere product-safety instrument (cfr., D. MÜGGE, *EU AI Sovereignty: for Whom, to What end, and How?*, in *Journal of European Public Policy*, 2024, p. 2200 ff.; S. POLI, *Il rafforzamento della sovranità tecnologica europea e il problema delle basi giuridiche*, in this *Journal*, 2021, p. 69 ff.; L. VON DITFURTH, *The European Union's Pursuit of Digital Sovereignty Through Legislation*, in *JIPITEC*, 2025, p. 286 ff.). It is a hybrid regulatory framework in which obligations for GPAI providers simultaneously serve the objectives of market standardisation and the protection of fundamental rights. The systemic-risk regime illustrates this hybridity by tying technical thresholds and governance mechanisms to rights-based constitutional principles. Indeed, systemic risks are not confined to product safety: they span democracy, discrimination, economy, environment, governance, security, and fundamental rights, among others, demonstrating how AI can generate harms that diffuse throughout social and economic systems (cfr., R. UUK ET AL., *A Taxonomy of Systemic Risks from General-Purpose AI*, in *RAND Working Paper Series*, 2024, p. 1 ff.). This hybrid logic finds concrete expression in other governance instruments of the AIA, most notably the Code.

4. Building on the discussion of GPAI models and systemic risk, the analysis now focuses on the associated obligations and on the role of the Code in supporting compliance with them.

The AIA establishes two main sets of obligations for providers of GPAI models. Art. 53 imposes baseline duties on all providers, including the requirement to maintain and update the technical documentation of the model (Art. 53(1) AIA) and to ensure compliance with applicable copyright rules (Art. 53(2) AIA). In addition to these duties, providers of GPAI models presenting systemic risk are required to: conduct and document state-of-the-art model testing to detect and mitigate systemic risks (Art. 55(1)(a) AIA); analyse and manage systemic risks at EU level arising from the development, marketing or use of such models (Art. 55(1)(b) AIA); track serious incidents and promptly report them, together with any corrective measure, to the AI Office and, where appropriate, national authorities (Art. 55(1)(c) AIA).

In both cases, compliance with the obligations set out in the AIA may be demonstrated through the use of codes of practice. This interplay is clarified in the AIA itself, which provides that: "The AI Office shall encourage and facilitate the drawing up of codes of practice at Union level in order to

contribute to the proper application of this Regulation, taking into account international approaches" (Art. 56(1) AIA). Under Art. 56(2) AIA, the AI Office, together with the European Artificial Intelligence Board, must ensure that the codes of practice reflect the requirements set out in Arts. 53 and 55. Moreover, the AIA expressly allows providers of GPAI models to rely on such codes to demonstrate compliance with their statutory obligations, pending the adoption of harmonised standards (Arts. 53(4) AIA and 55(2) AIA) and identifies codes of practice as a "central tool" for compliance (Recital No. 117 AIA). This gives the Code a distinctive, *quasi-binding* status (K. TAFOUR, *EU AI Obligations for GPAI Providers: Compliance, Enforcement & Deadlines (2025–2027)*, in *MediaLaws*, 28 August 2025). Thus, the Code, although not legally binding, provides that missing layer of clarity (T. MARCELIN, V. APRAXINE, *The European Union's AI Code of Practice*, in *EPRS - Digital Issues in Focus*, 2025) playing a central role within the EU's regulatory framework, which often combines hard law with softer instruments that shape its application.

While joining the Code remains optional, companies that adhere to it gain an accepted path benefit from having an officially recognized way to demonstrate that they comply with the corresponding AIA obligations, as in this case the supervisory scrutiny is largely directed at whether the provider adheres to the provisions of the Code, thereby increasing predictability and helping to limit administrative overhead (see here; see also L. SODER, E. PROVIĆ, *From Code to Compliance: The EU's Bid to Regulate General-Purpose AI*, in *Interface*, 17 July 2025). Those that choose not to participate must instead demonstrate conformity through other, often more burdensome, forms of documentation. Importantly, following the Code does not exempt providers from possible enforcement action, nor does it formally constrain the Commission's interpretation of the AIA.

The content of the Code is divided into three chapters, each setting out specific commitments and describing the measures that providers may implement in order to give practical effect to those commitments.

The first, on transparency, and the second, on copyright, apply to all providers GPAI models. In particular, the transparency chapter sets out how providers of such models must document, update, and share relevant information with competent authorities and downstream developers, ensuring the quality, integrity, and security of the data. On the other hand, the copyright chapter which sets out how providers must comply with EU copyright law in data collection and model use. It establishes measures to adopt copyright policies, avoid unlawful use of protected content, respect rights reservations, mitigate the risk of copyright-infringing outputs, and handle complaints from rightsholders.

The third chapter, on safety and security, is limited to providers of GPAI models presenting systemic risks. More specifically, it establishes frameworks for identifying, assessing, and mitigating systemic risks across the model lifecycle, setting obligations on safety measures, cybersecurity protections, incident reporting, and accountability structures. As to this chapter, early

drafts of the Code envisaged the inclusion of a taxonomy of systemic risks designed to provide more detailed criteria for identification (G. TORCHIO, S. GOODGER, *Regulating General Purpose AI in the EU: A Analysis of the AI Code of Practice*, in *European Policy Centre*, 23 September 2025, p. 2 ff.). More specifically, in the third draft of the Code, systemic risks were divided into "selected" and "other" types. Fundamental rights appeared only in the latter category, creating the impression of a hierarchy that placed them in a secondary position. This approach was criticized as inconsistent with the AIA and with the broader framework of EU law.

The final version adopts a different structure. It identifies broad categories of risk which include public health, safety, public security, fundamental rights and society at large. What had previously been labelled as "selected risks" are now reformulated as "specified systemic risks", covering CBRN threats, loss of control, cyber offence and harmful manipulation. The revision avoids an explicit hierarchy but shifts the emphasis toward a more securitized understanding of systemic risks.

5. What makes the Code particularly notable is the way it was created. Rather than being imposed through a traditional regulatory approach grounded in the ordinary legislative process, the Code was drafted through a collaborative process in which the very actors who will have to comply with it were involved in shaping its content (A. DATTA, *The Regulation of GPAI Model Providers Under the EU AI Act*, in M. DENGA, L. HORNUF (eds.), *Regulatory Competition on the Digital Economy*, Cham, 2025, p. 119 ff.). In less than a year, a process led by thirteen scientists and enriched by contributions from more than one thousand participants across industry, academia, civil society, and governments produced a detailed framework. The Code now serves as a practical reference for providers of GPAI models and as an early signal of how the Commission intends to apply the AIA. This participatory development also explains why the Code presents itself less as a rigid regulatory framework and more as a negotiated governance instrument.

In this sense, the Code exemplifies a broader recent tendency in EU digital regulation to complement hard law with co-regulatory instruments, as also exemplified by the DSA (E. COCCHIARA, *La regolazione del digitale: tra self-regulation, codici di condotta ed authorities*, in *Nuove Autonomie*, No. 1, 2025, p. 385 ff.; R. GRIFFIN, *Codes of Conduct in the Digital Services Act: Functions, Benefits & Concerns*, in *Technology and Regulation*, 2024, p. 167 ff.). This approach responds to the rapid evolution of technology, which risks rendering legislation quickly outdated. It is therefore more practical to entrust specifications to codes, as they can be updated more rapidly than hard law (R. GRIFFIN, *Codes of Conduct in EU Digital Regulation and AI Policy: The Potential and Risks of Soft Law Tools*, in K. PRIFTI ET AL. (eds.), *Digital Governance: Confronting the Challenges Posed by Artificial Intelligence*, The Hague, 2024, p. 253 ff.).

However, the use of soft law instruments becomes particularly sensitive when they operate in regulatory domains with direct implications for

fundamental rights, democratic values, and systemic risks with implications for society as a whole. Yet their development – as illustrated by the case of the Code of Practice for GPAI models – takes place through informal negotiations between unelected authorities and selected stakeholders, that often may be moved by industrial interests rather than concerns on the fundamental rights of citizens. This institutional setting raises legitimacy concerns when technical choices embedded in the Code affect rights already recognized under the EU legal framework. Under the DSA framework, for example, codes of practice contain measures to address illegal content to mitigate systemic risks have direct implications for fundamental rights protected by the Charter, including freedom of expression and information (Art. 11 of the Charter), non-discrimination (Art. 21, also covering equal access), and the protection of human dignity (Art. 1). Similarly, under the AIA, the Code of practice for GPAI models establishes obligations concerning documentation, transparency, and systemic risk governance, all of which impact fundamental rights such as privacy and data protection (Arts. 7 and 8), intellectual property (Art. 17(2)) as well as the right to an effective remedy (Art. 47). These rights are also referred to in Recital No. 48 AIA, which identifies the extent of adverse impacts on fundamental rights as a key criterion for the classification of AI systems as high-risk. Indeed, in these contexts, the choice of metrics, technical parameters, and operational procedures is never neutral: it directly shapes the level of guarantees offered to end-users and determines the actual effectiveness of the protection of their rights.

Against this backdrop, the Code can be viewed as an instrument that seeks to operationalize one of the core tensions embedded in EU digital law: the balance between innovation and the protection of fundamental rights (cfr., A. BRADFORD, *The False Choice Between Digital Regulation and Innovation*, in *Northwestern University Law Review*, 2024, p. 377 ff.; M. GILLE, M. TROMPMANN-FRICK, T. SCHOMACKER, *Balancing Public Interest, Fundamental Rights, and Innovation: The EU's Governance Model for Non-High-Risk AI Systems*, in *Internet Policy Review*, 2024, p. 2 ff.). Indeed, the Code explicitly recalls Article 1(1) AIA, which combines a high level of protection for health, safety, and fundamental rights with the fostering of AI development and uptake within the Union. This dual objective makes the Code a testing ground for the EU's broader regulatory philosophy in the AI domain, characterised by a human-centric approach to AI (cfr., Y. S. BALCIOĞLU, A. A. ÇELIK, E. ALTINDAĞ, *A Turning Point in AI: Europe's Human-Centric Approach to Technology Regulation*, in *Journal of Responsible Technology*, 2025, p. 1 ff.; A. PIROZZOLI, *The Human-Centric Perspective in the Regulation of Artificial Intelligence*, in *europeanpapers.eu*, 2024, p. 105 ff.) in which innovation is consistently embedded within a framework that prioritize the protection of fundamental rights (A. BRADFORD, *op. cit.*, p. 2 ff). If obligations were excessively strict, they could impair the effectiveness of the internal market and weaken the EU's global competitiveness. Conversely, if protections were too lax, they would undermine the substance of

fundamental rights, eroding the legitimacy of innovation itself (G. TORCHIO, S. GOODGER, *op. cit.*, p. 3).

The willingness to preserve this delicate equilibrium within the Code is clearly reflected in provisions that embody a dynamic risk-based logic: compliance obligations remain light when risks are limited, but progressively tighten as risks emerge and escalate. For example, in the Safety and Security Chapter, the Code requires that systemic-risk assessments be conducted iteratively along the entire model lifecycle, but "proportionate to the risks", avoiding duplication where previous assessments remain valid, and allowing "lighter-touch model evaluations […] at appropriate trigger points" when risks remain limited. This graduated approach prevents compliance obligations from becoming excessively rigid: innovation can proceed at speed and providers retain the flexibility to update models in response to technological progress, thereby sustaining internal market effectiveness and the EU's global competitiveness. At the same time, when risks increase, the Code mandates escalation to a full systemic-risk assessment, ensuring responsive protection of fundamental rights that may be affected by emergent model capabilities.

Finally, it is essential to emphasize that the Code does not bind the Commission's interpretation of the AIA, nor does it limit the possibility of future enforcement actions or regulatory corrections. Precisely because of its non-binding and flexible character, the Code preserves significant room for adjustment: it may be amended or supplemented as new evidence or societal concerns emerge during implementation (C. GALLESE, *The GPAI Code of Practice, a Long Journey not Over Yet*, in *MediaLaws*, 8 September 2025). This openness ensures that safeguards for fundamental rights can be continuously reinforced in step with technological developments, rather than being locked into static solutions that may soon become inadequate. This adaptive logic directly reflects the very objective of the AIA to reconcile innovation and economic freedoms with the effective protection of privacy, non-discrimination, and human dignity. It is therefore the Code's ability to evolve – rather than to crystallize technical choices – that allows regulatory requirements to remain proportionate and context-sensitive, preventing the emergence of horizontal conflicts between providers, users, and supervisory authorities within the internal market.

6. The EU framework for GPAI models shows that the AIA does not simply extend product safety logic to emerging technologies. It introduces a hybrid regulatory paradigm in which the protection of fundamental rights and other EU objectives become closely connected. The Code of Practice for GPAI models is a concrete expression of this ambition. By filling the gap between legal obligations and technical compliance practices, the Code operates as an early implementation tool of the AIA, complementing and specifying its provisions while at the same time preserving room for adaptation in the face of fast-changing technological developments. Its co-regulatory nature however rests on a delicate balance. The speed and efficiency offered come at the price of relying on informal negotiations that

can shape decisions with direct implications for rights protected by the Charter.

On the one hand, the flexibility of the Code helps avoiding a regulatory framework that would slow innovation, weaken technological sovereignty, and harm the Union's global competitiveness. However, this same flexibility calls for strong guarantees of democratic oversight, transparency, and accountability that draw on existing EU frameworks, such as the European Declaration on Digital Rights and Principles, in order to inform an interpretation of the Code of Practice that is firmly oriented towards the EU's values and fundamental rights. Without these safeguards there is a risk of an implicit transfer of fundamental rights protection to private actors. This makes it especially important for the AI Office and national authorities to take an active role and ensure that the technical choices embedded in the Code do not reduce the level of protection granted to end users.

The Code should therefore be seen not as a final point but as a dynamic instrument. It must evolve in line with empirical evidence, scientific progress and the emergence of new forms of systemic risk. Its adaptive character is not a weakness but a defining element of the EU approach. Only a regulatory framework that remains proportionate, contextual and open to revision can keep pace with technological innovation while still safeguarding the essential core of fundamental rights.

In conclusion, the attempt of the EU to "tame systemic risk" through a co-regulatory approach to GPAI confirms that the balance between innovation and rights is not a contradiction. It is an ongoing process. The success of the Code and eventually of the AIA will depend on the ability of EU institutions to ensure that flexibility serves the public interest. At the same time the framework must remain agile enough to support an environment where technology is not only safe but also continues to grow inside the EU.

ABSTRACT (ITA)

L'ascesa dei modelli GPAI ha messo in discussione il quadro originario in materia di gestione del rischio dell'*Artificial Intelligence Act* (regolamento (UE) 2024/1689, AIA). Accanto a una nuova categoria giuridica è emersa l'esigenza di considerare regole sul rischio sistemico. Questo post analizza il Codice di condotta GPAI dell'UE come strumento fondamentale nell'attuazione dell'AIA, traducendo obblighi generali in pratiche concrete in materia di trasparenza, proprietà intellettuale e sicurezza, distinguendo tra tutti i fornitori di GPAI e quelli i cui modelli presentano un rischio sistemico. Sebbene formalmente volontario, il Codice ha una forza quasi vincolante in quanto imposta traiettorie rilevanti verso la conformità e la certezza giuridica, nel tentativo di bilanciare tensioni tra innovazione e tutela dei diritti. A tale proposito, si sostiene che il successo del Codice dipenderà dalla sua capacità di rimanere sufficientemente flessibile da sostenere la competitività e l'innovazione nell'UE, garantendo al contempo che la tutela dei diritti fondamentali non sia delegata ad attori privati.

ABSTRACT (ENG)

The rise of GPAI models has challenged the original risk-based framework of the Artificial Intelligence Act (Regulation (EU) 2024/1689, AIA), requiring a new legal category and the introduction of systemic risk as a regulatory cornerstone. This post analyses the EU's GPAI Code of Practice as a pivotal instrument in the implementation of the AI Act, translating broad obligations into concrete practices on transparency, copyright, and safety, distinguishing between all GPAI providers and those whose models present systemic risk. Although formally voluntary, it carries *quasi-binding* force by offering a recognised pathway to compliance and legal certainty, while revealing tensions between innovation and rights protection. In sum, the article argues that the success of the Code will depend on its ability to remain flexible enough to support competitiveness and innovation in the EU while ensuring that the protection of fundamental rights is not delegated to private actors.