

# The AI Act's GPAI Code: Hidden Policy Choices

---

By Cornelia Kutterer and Theodoros Karathanasis



## The AI Act's GPAI Code: Hidden Policy Choices

**This article examines the General-Purpose AI (GPAI) Code of Practice (CoP), a critical component of the EU's AI regulation framework. It analyzes specific textual ambiguities and interpretive flexibilities within the legal provisions that grant the AI Office substantial discretionary authority in implementation. By identifying these strategic openings in the regulatory text, the article illuminates how different interpretations could significantly shape enforcement priorities and compliance requirements.**

On December 19<sup>th</sup>, 2024, the European Commission (EC) introduced the [second draft](#) of the Code of Practice (CoP) for General-Purpose AI (GPAI),<sup>1</sup> an important addition to its regulatory framework for artificial intelligence (AI). However, the ongoing debates on the GPAI Code are overshadowed by various stakeholders' conflicting interests and increasing political pressure.

Supplementing the European Union's legal framework on AI ([EU AI Act](#)) that went into effect August 1, 2024, this draft Code aims to detail the rules for [providers of GPAI models](#) and [GPAI models with systemic risks](#).<sup>2</sup> Although the third draft of the CoP was expected for the week of February 17, 2025, it has been delayed by at least one month, according to a new European Commission [timeline](#), because the CoP working group chairs want to ensure that the latest document [reflects](#) the feedback they received from stakeholders, as well as to make it legally robust. Indeed, it appears that some industry players executives have [criticized](#) the draft rules.

The article **contextualizes the GPAI CoP** within the EU's broader approach to tech regulation, highlighting the **role of co- and self-regulation (1.)**. Furthermore, future **AI Office policy decisions** are identified as critical in **shaping the GPAI CoP's path forward**, particularly in addressing challenges related to stakeholder interests and potential over-regulation (2.). As a preliminary proposal, a **modular**

**approach** is put forward to further streamline discussions and enhance clarity. This approach advocates for **separate codes** for intellectual property, transparency, and risk mitigation (3.). As a secondary proposal, the CoP should be framed and utilized as a **living framework (4.)**. This would **facilitate compliance**, support the development of best practices, and enable voluntary enhancement in AI governance.

### 1. Putting the GPAI CoP into context: The role of Co- and Self-Regulation in tech policies

The EU's approach to regulating digital technologies has undergone a remarkable transformation over the past decade when it comes to soft law and co-regulation.<sup>3</sup> Initially resistant to soft law instruments like codes of practice - viewing them as insufficiently robust for European regulatory traditions - EU policymakers have gradually embraced them as essential tools for governing complex technological systems, and addressing at EU level topics that were not ripe for EU regulation.

The EU's journey with self-regulation began tentatively. Defined as a "large number of practices, common rules, codes of conduct (CoC) and voluntary agreements with economic actors, social players,

<sup>1</sup> GPAI model refers to "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" ([Article 3](#), point 63 of the AI Act).

<sup>2</sup> [Article 56](#) of the AI Act (AIA).

<sup>3</sup> Terpan, F. (2015), Soft Law in the European Union. *European Law Journal*, 21: 68-96. <https://doi.org/10.1111/eulj.12090>

NGOs and organized groups (...).<sup>4</sup> Self-regulation was often dismissed as inadequate substitutes for hard law in its early attempts at industry.<sup>5</sup> The UK's longer tradition with codes of conduct stood in stark contrast to Continental scepticism. However, the integration of a regulatory backstop in the unfair commercial practice directive enabled and the rapidly evolving nature of digital technologies forced a rethinking of this approach. While the Treaty on the Functioning of the European Union (TFEU) provides the general legal foundation for co- and self-regulatory instruments, each code developed at the EU level has evolved with distinct features tailored to its specific sectoral challenges and governance objectives.

This evolution toward co-regulation is evidenced through several significant milestones. The [2016 Code of Conduct on Illegal Hate Speech](#) marked an early experiment, bringing major platforms together to address content moderation challenges. The [2018 Code of Practice on Disinformation](#) represented a more ambitious effort, though it initially faced criticism for lacking enforcement mechanisms. The strengthened [2022 Code of Practice on Disinformation](#) later demonstrated the potential for codes to evolve into more robust instruments. Finally, the [Digital Services Act's \(DSA\)](#) incorporation of codes as co-regulatory tools signaled their mainstream acceptance in EU digital policy.

The [GDPR's Code of Conduct framework](#), serving primarily as a compliance mechanism, represents perhaps the closest precedent for the AI Act's forthcoming CoP for GPAI models. More than the CoP on disinformation, GDPR CoCs serve as a **blueprint for structured co-regulation**, linking voluntary commitments to regulatory oversight. These codes provide a formalized mechanism for demonstrating compliance with GDPR obligations and, in some cases, can serve as a basis for cross-border data transfers under the regulation. The AI

Act's CoP mirrors this regulatory architecture, creating an industry-driven yet enforceable framework to guide compliance with GPAI obligations.

Yet, regarding enforcement authority involvement and stakeholder participation, the AI Act's CoP incorporates critical lessons from previous co-regulatory experiences, establishing a more precisely defined and substantive role for the Commission and other stakeholders. The [AI Office](#) plays a central role in the development of this CoP through both encouraging and facilitating their drafting, review and adaptation.<sup>6</sup> It is the responsibility of the AI Office and AI Board to ensure that CoP, at a minimum, addresses the obligations in Articles [53](#) and [55](#) of the AI Act.<sup>7</sup> This could include among others ensuring that information is kept up to date with market and technological changes, providing sufficient detail on training data summaries,<sup>8</sup> identifying the type, nature, and sources of systemic risks in GPAI models,<sup>9</sup> or even outlining measures for assessing and managing these risks, with documentation as needed.<sup>10</sup>

The code is uniquely embedded in the GPAI model provision enforcement structure, as the Commission explicitly considers "commitments made in relevant codes of practice in accordance with Article 56" when determining fines of up to 3% of annual worldwide turnover or €15 million for providers who intentionally or negligently infringe the regulation, fail to comply with information requests, disregard Commission measures, or deny access for model evaluations—establishing a direct link between code adherence and potential penalty mitigation.<sup>11</sup>

<sup>4</sup> Communication from the Commission, Action plan 'Simplifying and improving the regulatory environment', COM(2002) 278 final, 5 June 2002, p. 11.

<sup>5</sup> The Federation of European Direct and Interactive Marketing ('FEDMA') Community Code of Conduct on direct marketing is one of the few examples of approved codes of conduct under Article 27 of the repealed Data Protection Directive (DPD) - Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with

regard to the processing of personal data and on the free movement of such data.

<sup>6</sup> Article 56(1) and Recital 116 of the AIA

<sup>7</sup> Article 56(2) of the AIA

<sup>8</sup> Article 56(2)(a)–(b) of the AIA

<sup>9</sup> Article 56(2)(c) of the AIA

<sup>10</sup> Article 56(2)(d) of the AIA

<sup>11</sup> Article 101.

## 2. AI Office Policy Decisions Critical for GPAI CoP's Path Forward

GPAI model providers and the AI Office can make deliberate decisions that will impact the compliance and enforceability of the code. A degree of regulatory positioning by the AI Office is thus expected, as it seeks to shape a framework that optimizes compliance incentives while navigating political pressures and industry dynamics.

While the draft CoP outlines detailed requirements for transparency, copyright compliance, and systemic risk mitigation, its adoption faces challenges from divergent stakeholder interests and potential political resistance to perceived over-regulation. Over-regulation can occur when statutory obligations overlap, leading to potential confusion and duplication of effort.<sup>12</sup> In highly innovative environments, openness to innovation may be hindered by legal uncertainty. To solve the conflict between legal principles and legal uncertainty, the regulator can add procedural instruments that enable the regulation addresses to increase legal certainty on their own by setting up standards by private entities.<sup>13</sup> A co-regulation strategy is useful for assessing risks caused by innovation because it enables the regulation addressees to consider the specific context, whether concerning a particular product, service category, or type of processing activity.<sup>14</sup>

Despite the unprecedented effort that is currently undertaken to finish the code in time, the development of the GPAI CoP faces an uncertain future, driven not only by intense lobbying efforts and new political tensions, but also due to textual ambiguities that create interpretive latitude and complex enforcement dynamics.

Certain provisions, in particular adequacy decisions, common rules, general validity and their relation to common standardization lack explicit definitional boundaries (**Table 1**).

<sup>12</sup> von Grafenstein, M. (2022). "Chapter 19: Co-regulation and competitive advantage in the GDPR: Data protection certification mechanisms, codes of conduct and data protection-by-design". In *Research Handbook on Privacy and Data Protection Law*. Cheltenham, UK: Edward Elgar Publishing.

<sup>13</sup> Grafenstein, Maximilian, *Co-Regulation and the Competitive Advantage in the GDPR: Data Protection Certification Mechanisms, Codes of Conduct and the 'State of the Art' of Data Protection-by-Design* (February 18, 2019). Forthcoming in González-Fuster, G., van Brakel, R. and P. De Hert *Research Handbook on Privacy and Data Protection Law*. Values, Norms

and Global Politics, Edward Elgar Publishing., Available at SSRN: <https://ssrn.com/abstract=3336990>

<sup>14</sup> See also Beckers, Anna, *Regulating Corporate Regulators Through Contract Law? The Case of Corporate Social Responsibility Codes of Conduct* (June 3, 2016). EUI Working Papers - Max Weber Programme 2016/12, Available at SSRN: <https://ssrn.com/abstract=2789360>; Vavrečka, Jan & Stepanek, Petr. (2014). *Self-regulation of Advertising – Controversial Impacts of Ethical Codes*. 10.13140/2.1.2579.424; Biedermann, R. (2006). From a weak letter of intent to prevalence: the toy industries' code of conduct. *Journal of Public Affairs*, 6, 197-209.

<b>Table 1. AI Act’s provisions on Adequacy Decisions, Common rules and General Validity in the context of the CoP for GPAI Models</b>	
	<p><b>Article 52</b></p> <p>4. Providers of general-purpose AI models may rely on codes of practice within the meaning of <a href="#">Article 56</a> to demonstrate compliance with the obligations set out in paragraph 1 of this Article, <b>until a harmonised standard is published</b>. Compliance with European harmonised standards grants providers the presumption of conformity to the extent that those standards cover those obligations. Providers of general-purpose AI models who do not adhere to an approved code of practice or do not comply with a European harmonised standard shall demonstrate alternative adequate means of compliance for assessment by the Commission.</p>
<p><b>Recital 117</b></p> <p>The codes of practice should represent a central tool for the proper compliance with the obligations provided for under this Regulation for providers of general-purpose AI models. Providers should be able to rely on codes of practice to demonstrate compliance with the obligations. <b><i>By means of implementing acts, the Commission may decide to approve a code of practice and give it a general validity within the Union, or, alternatively, to provide common rules for the implementation of the relevant obligations, if, by the time this Regulation becomes applicable, a code of practice cannot be finalised or is not deemed adequate by the AI Office.</i></b> Once a harmonised standard is published and assessed as suitable to cover the relevant obligations by the AI Office, compliance with a European harmonised standard should grant providers the presumption of conformity. Providers of general-purpose AI models should furthermore be able to demonstrate compliance using alternative adequate means, if codes of practice or harmonised standards are not available, or they choose not to rely on those.</p>	<p><b>Article 56</b></p> <p>1. 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.</p> <p>6. The AI Office and the Board shall regularly monitor and evaluate the achievement of the objectives of the codes of practice by the participants and their contribution to the proper application of this Regulation. The AI Office and the Board shall assess whether the codes of practice cover the obligations provided for in Articles 53 and 55, and shall regularly monitor and evaluate the achievement of their objectives. <b><i>They shall publish their assessment of the adequacy of the codes of practice. The Commission may, by way of an implementing act, approve a code of practice and give it a general validity within the Union.</i></b> That implementing act shall be adopted in accordance with the examination procedure referred to in Article 98(2).</p> <p>9. Codes of practice shall be ready at the latest by 2 May 2025. The AI Office shall take the necessary steps, including inviting providers pursuant to paragraph 7. <b><i>If, by 2 August 2025, a code of practice cannot be finalised, or if the AI Office deems it is not adequate following its assessment under paragraph 6 of this Article, the Commission may provide, by means of implementing acts, common rules for the implementation of the obligations provided for in Articles 53 and 55, including the issues set out in paragraph 2 of this Article.</i></b> Those implementing acts shall be adopted in accordance with the examination procedure referred to in Article 98(2).</p>

Source: Eur-Lex

When the AI Office deems a code adequate (publishing this determination as an official adequacy decision), signatories can use their adherence as a presumption of conformity with regulatory requirements. Non-signatories, however, must independently demonstrate equivalent safeguards through alternative compliance methods, potentially subjecting them to heightened scrutiny from regulators. If the AI Office determines a code to be inadequate, the Commission may—but is not obligated to—develop and adopt common rules that would exist alongside the industry code.<sup>15</sup> These Commission-created rules would serve as authoritative guidance for all providers in implementing their regulatory obligations, regardless of whether they had previously signed the industry-developed code.<sup>16</sup> This discretionary power to adopt or withhold common rules represents yet another significant level of influence for the Commission, granting it substantial control over the compliance trajectory.

Another critical interpretive ambiguity, or twist, lies in the "general validity" clause. General validity could suggest a move beyond the immediate participants in the co-regulatory effort and align with the broader applicability of common rules, but the precise meaning and scope of general validity is not defined within the AI Act itself.<sup>17</sup>

Both the AI Act and the GDPR empower the European Commission to grant general validity to codes of conduct or practice within the EU.<sup>18</sup> In both frameworks, the codes are seen as instruments to aid in the application and specification of the respective regulations, serving as a compliance tool.<sup>19</sup> However, there are key differences between the GDPR and the AI Act, most importantly the fact that its enforcement is, in the case of GDPR,

undertaken by independent national DPAs.<sup>20</sup> However, these established rationales for formal recognition in the GDPR are not essential prerequisites for the AI Act's CoP given the exclusive powers to supervise and enforce Chapter V of the AI Act.<sup>21</sup> Even in scenarios where national regulators evaluate compliance during enforcement actions against high-risk AI systems, a regulatory interaction that involves the AI Office and falls within the Commission's supervisory domain, the Commission's adequacy declaration would likely provide sufficient legal standing for the CoP, establishing an enforcement framework that diverges significantly from the GDPR code of conduct framework. The development of CoP for GPAI models under the EU AI Act could mirror Germany's *general validity mechanism*<sup>22</sup> in labour law where sectoral agreements gain binding force for all market players through state endorsement to prevent competitive distortions. If the CoP is formally adopted and deemed "generally valid" via implementing act, under this interpretation, it could function similarly – transforming voluntary guidelines into enforceable compliance standards for *all* GPAI providers, regardless of their participation in the CoP drafting process to avoid that those participating in the code are disadvantaged. This interpretation would ensure coherence with the common rules applicable to all providers in cases the Commission has decided the CoP is inadequate. More importantly, this interpretation would enable the AI Office to monitor market practices of non-participating providers and empower it to impose these rules universally, thereby capturing the complete trajectory of governance mechanisms and their progressive maturation across the entire GPAI ecosystem.

<sup>15</sup> Article 56(9) states that codes of practice should be ready at the latest by 2 May 2025. However, if a code of practice has not been finalized or is deemed inadequate by 2 August 2025 (Article 56(9)), it is up to the Commission to provide guidance on the obligations laid down in Articles 53 and 55 as well as Article 56(2) (Article 56(9)). The existence of conflicting interests among the various stakeholders regarding the AI Office approach toward a CoP for GPAI could however jeopardize its adoption before May 2025.

<sup>16</sup>The code could additionally be enforced under [Article 6\(2\)\(b\) of the Unfair Commercial Practice Directive \(UCPD\)](#) as the linchpin that enforces adherence to self-imposed codes of conduct.

<sup>17</sup> As of now, the Commission has not yet provided more guidance but told the authors they would - following a request for clarification.

<sup>18</sup> Article 56 of the AIA and Article 40 of the GDPR

<sup>19</sup> The GDPR CoC general validity, which has not been used as of today, was - again - the possible blueprint for including this measure in the AI Act.

<sup>20</sup> The GDPR makes the submission procedure obligatory for EU-wide codes, requiring a national authority to elect co-reviewing supervisory authorities

<sup>21</sup> Article 88.

<sup>22</sup><https://www.dbb.de/lexikon/themenartikel/a/allgemeinverbindlicherklaerung.html>

The Commission's ability to grant general validity to the CoP represents a powerful regulatory lever, particularly against strategic non-signatories who calculate that enforcement resources will initially target those who formally committed but failed to comply—overlooking that signatories may face additional financial penalties for substantive violations of the underlying CoP commitments. If major GPAI developers decline to sign the CoP, this could create an uneven playing field, where those who voluntarily commit to higher standards face greater scrutiny than those who choose not to opt in.

The Commission could of course reduce the scope of general validity to jurisdictional effect, similar to that of the GDPR, but such a limitation appears logically inconsistent with the regulatory architecture outlined above. Both interpretations of general validity present defensible positions, and so far, the Commission has postponed taking a definitive stance, even though this ambiguity could significantly influence providers' willingness to participate in codes that are nominally intended to remain voluntary. The AI Office will eventually need to make a consequential policy determination that takes these issues into account while balancing regulatory objectives against legitimate stakeholder concerns regarding compliance efficiencies and procedural fairness.

The last cliffhanger is the Damocles sword hanging over the CoP. Article 52.4 states that "Providers of general-purpose AI models may rely on codes of practice within the meaning of Article 56 to demonstrate compliance with the obligations set out in paragraph 1 of this Article, until a harmonised standard is published." This provision reveals that the entire co-regulatory effort could ultimately be superseded by formal technical standards, should the Commission decide to initiate a standardization request under EU law for GPAI models, and should [CEN-CENELEC](#) successfully develop and deliver such standards.<sup>23</sup>

<sup>23</sup> Although the Commission adopted an implementing decision on 22 May 2023 on a standardisation request to CEN-CENELEC in support of the Union's policy on artificial intelligence ([C\(2023\)3215](#)), it did not include GPAI models.

The Commission must carefully assess these perks and cliffhangers to develop a broader and clearer interpretive framework. A robust and regularly updated CoP may prove more progressive than a standardization framework that struggles to accommodate rapidly evolving considerations. Ultimately, the future of the GPAI CoP hinges on these political determinations. While it holds significant potential to become a flexible and widely accepted governance tool, its precise role within the broader AI regulatory framework remains uncertain—subject to both legal interpretation and political will.

### 3. Modular Approach: A Case for Separate Codes

The development of the CoP for GPAI involves multiple stakeholder groups, each with distinct priorities and, in some cases, conflicting interests. The level of tension between these groups varies depending on the issue at hand. Recognizing this complexity, the AI Office has already taken steps to manage these divergent perspectives by forming separate working groups during the drafting process.<sup>24</sup> However, while these groups address specific topics, they remain part of a single, overarching CoP rather than distinct modules. Indeed, a modular approach to the CoP would offer a number of important advantages that would make the regulatory framework more effective and more adaptable.

First, it facilitates targeted engagement, allowing stakeholders to focus on specific areas of concern without being required to negotiate all aspects of a monolithic code. This ensures that discussions remain relevant to each group's expertise and interests, reducing unnecessary conflicts and delays. Second, it accelerates consensus-building, as less controversial modules can be finalized and implemented more quickly, while more complex areas—such as risk mitigation—can undergo further refinement without holding back the entire CoP. This staged approach enables regulatory progress without requiring immediate agreement on every

<sup>24</sup> This process is structured around four working groups: Transparency and Copyright, Categorization of Risk and Assessment, Identification of Mitigation Measures, and Governance and Internal Risks Assessments.

detail. Third, it enhances adaptability, since AI technologies and risks continue to evolve. By structuring the CoP into distinct modules, individual sections can be updated or revised independently, avoiding the need for a complete regulatory overhaul each time an adjustment is required. Fourth, it provides compliance flexibility, allowing AI providers to demonstrate adherence through different combinations of modules. This approach recognizes that not all providers face the same regulatory challenges, making it easier for them to align their compliance efforts with their operational models while still meeting EU requirements. Overall, a modular CoP ensures that the AI Act's objectives are met while promoting a more manageable, efficient, and future-proof regulatory framework.

Given the nature of stakeholder engagement, a more structured, modular approach—dividing the CoP into Intellectual Property, Transparency, and Risk Mitigation—would therefore further streamline discussions and enhance regulatory clarity. [Article 56.7](#) of the AI Act further supports this approach by explicitly referencing multiple codes and permitting providers of non-systemic-risk GPAI models to adhere only to the obligations outlined in Article 53. This reinforces the rationale for structuring the Code of Practice as separate, modular codes, allowing for differentiated compliance pathways that better align obligations with the varying risk profiles and capacities of AI providers.

A dedicated Intellectual Property Code is justified due to the legal and technical intricacies surrounding copyright compliance in AI training.<sup>25</sup> The use of copyrighted material in AI development—particularly concerning [Text and Data Mining \(TDM\)](#) exceptions—has already been identified in the AI Act as requiring careful balancing between innovation and rights protection.<sup>26</sup> A separate module could provide specific, enforceable guidance on how AI developers should engage with rightsholders, implement licensing frameworks, and resolve

disputes over copyrighted content. This would also align with the EU Commission's January 28, 2022, initiative to establish a [CoP for the smart use of Intellectual Property](#), ensuring coherence across AI and copyright governance.

A standalone Transparency Code is necessary because transparency requirements touch on multiple areas of contention, particularly between AI providers, civil society groups, and rights holders. GPAI developers often resist broad disclosure obligations, whereas other stakeholders demand comprehensive transparency measures, particularly regarding training data sources, model architecture, and system limitations. Transparency is a cornerstone of trust and accountability, making it essential to develop clear documentation standards that facilitate compliance, auditing, and oversight. Additionally, transparency about data provenance is critical for assessing privacy and data protection risks, further justifying the need for targeted, enforceable guidelines within a distinct module.

Risk mitigation strategies for systemic-risk GPAI models are another area where [stakeholder interests](#) diverge sharply. AI providers often prefer self-assessment frameworks, whereas regulators and civil society groups advocate for mandatory third-party audits—particularly for high-risk applications. A separate Risk Mitigation Code would allow for a more nuanced, risk-tiered approach, where independent audits are required for high-risk models, while lower-risk systems could rely on self-assessment with regulatory oversight. Given the evolving understanding of AI risks, such a framework would provide the necessary flexibility to adapt to emerging threats without imposing undue burdens on innovation.

<sup>25</sup> Kaigeng Li, Hong Wu, Yupeng Dong, (2024) 'Copyright protection during the training stage of generative AI: Industry-oriented U.S. law, rights-oriented EU law, and fair remuneration rights for generative AI training under the UN's international governance regime for AI', *Computer Law & Security Review*, Vol.55, <https://doi.org/10.1016/j.clsr.2024.106056>

<sup>26</sup> See also Theodoros Karathanasis, EU Copyright Directive: A 'Nightmare' for Generative AI Researchers and Developers? AI

Regulation Papers 23-10-2, AI-Regulation.com, October 17th, 2023, <https://ai-regulation.com/eu-copyright-directive-a-nightmare-for-gai/>; Eleni Polymenopoulou, Artificial Intelligence and The Future of Art: The Challenges Surrounding Copyright Law And Regulatory Action, AI Regulation Papers 23-10-1, AI-Regulation.com, October 6th, 2023, <https://ai-regulation.com/artificial-intelligence-and-the-future-of-art/>



#### 4. The Code as a Living Framework

The ongoing debate about the AI Act's CoP reflects a fundamental misunderstanding of the role and potential of such instruments. Industry stakeholders argue that the CoP should not impose obligations beyond the AI Act's explicit requirements, while civil society organizations push for stronger measures, such as mandatory external assessments. This framing, however, creates a false dichotomy — one that overlooks the multi-faceted purpose of codes of practice within the EU regulatory framework.

Like other EU digital regulations, the AI Act envisions codes serving three distinct functions. First, they facilitate compliance by translating legal obligations into technical specifications, establishing common interpretations of requirements, and creating standardized compliance frameworks. Second, they support best practice development, helping industries identify emerging standards, document proven risk mitigation approaches, and share successful implementation strategies. Third, they provide a mechanism for voluntary enhancement, allowing companies to commit to additional safeguards, take leadership roles in responsible AI development, and set aspirational benchmarks for ethical and transparent AI.

The current discourse on the CoP underestimates its strategic value by focusing solely on whether it should impose additional regulatory burdens. While the CoP remains a voluntary instrument—meaning companies may choose alternative means of demonstrating compliance—it has the potential to establish industry-wide best practices that go beyond the AI Act's baseline requirements.

The voluntary nature of the Code of Practice can operate on two distinct levels that require separate consideration. First, there's the voluntary adoption of the code itself—a provider's choice to become a signatory rather than complying through alternative means. Second, there are voluntary commitments within the code that exceed minimum regulatory requirements once a provider becomes a signatory. The challenge is maintaining this distinction; if "voluntary" commitments face the same enforcement as mandatory requirements, providers lose incentive to exceed minimums. To maintain the

voluntary nature of these commitments and prevent them from being absorbed into the enforcement structure, several approaches could be considered. First, the Commission could establish a clear firewall between voluntary CoP commitments and mandatory regulatory requirements, explicitly stating that additional voluntary measures undertaken through the CoP will not trigger enforcement actions beyond the baseline obligations. Second, implementing a "safe harbour" provision could protect organizations that demonstrate good-faith efforts to comply with their CoP commitments from punitive enforcement. Third, the Commission could develop a tiered approach that distinguishes between core compliance requirements and aspirational best practices within the CoP framework.

Many leading AI companies already exceed minimum regulatory expectations by implementing enhanced model documentation, robust testing protocols, and voluntary external audits. These self-imposed measures not only demonstrate the feasibility of higher standards but also highlight the industry's capacity for self-regulation and innovation in responsible AI.

A well-structured CoP can formalize these emerging best practices while preserving flexibility for companies to engage with them voluntarily. This approach respects corporate autonomy while setting clear industry benchmarks for excellence in AI governance. Moreover, it creates a structured pathway for organizations seeking to differentiate themselves as leaders in responsible AI development.

By recognizing the CoP as more than just a compliance tool—but rather as a framework for industry progress—the debate can shift from regulatory constraints to opportunities for proactive AI governance. However, this reframing, while effective for elements clearly beyond the AI Act's scope, provides little resolution for contentious interpretation issues regarding the Act's actual requirements—as evidenced by ongoing disputes about whether external assessments fall within the regulation's mandatory obligations or represent an extension beyond its text.

## Conclusion

The future of the GPAI CoP depends on balancing legal certainty with critical policy decisions by the AI Office during the final negotiations and its future enforcement. As a co-regulatory instrument bridging private and public governance, the interpretation and definition of the AI Office can significantly alter the trajectory of the enforcement tool.

The Commission's power to extend the scope of the code through an implementing act represents a potent regulatory tool, particularly for addressing strategic non-signatories, and prevent market distortions by establishing universal technical baseline requirements.

The CoP's effectiveness ultimately hinges on political decisions, industry alignment, and regulatory evolution. Given the diverse stakeholders with competing priorities, a modular approach separating Intellectual Property, Transparency, and Risk Mitigation could enhance clarity and streamline negotiations. The CoP should function primarily as an implementation mechanism facilitating compliance while supporting innovation in governance practices.

Industry stakeholders have cautioned against exceeding the AI Act's explicit requirements, while civil society advocates for stronger measures. The AI Office must balance effective regulation with contextual flexibility while building provider confidence in the CoP as a reliable compliance mechanism.

Should the Office fail to finalize or approve an adequate CoP, the Commission may establish common rules, potentially creating compliance challenges for providers navigating complex implementation requirements.



Cornelia Kutterer is a Research Fellow at the Chair. Her research interest is the interplay of technology law and advances in fairness, accountability, and transparency of socio-technical systems as well as applied AI governance and ethics in organizations. She is a qualified German lawyer and holds a law degree from the University of Hamburg and a master's degree of information technology and telecommunication law from the Strathclyde university in Glasgow. She worked for 15 years at the forefront of digital innovation and responsible policies at Microsoft, including on responsible AI, privacy, human rights, safety, law enforcement, content, and market regulation. Before joining Microsoft, she led the legal department of the European Consumer Organisation BEUC and gained experience in a top 10 law firm and the European Parliament.



Dr. Theodoros Karathanasis holds a PhD in European law from the Faculty of Law of the University of Grenoble Alpes. He is attached to the Centre for International Security and European Cooperation Studies (CESICE) and his research work in the field of cybersecurity has been funded by the Grenoble Alpes Institute of Cybersecurity. He is a member of the cyber experts network of the European Centre of Excellence for Combating Hybrid Threats (HybridCoE), as well as the EU CyberNet expert network. He is a Research Fellow at the AI-Regulation.com Chair.

**Cover Photo:** Created by **DALL-E** and **Adobe Photoshop** prompt by **Shadée Pinto**, Research Fellow at the Chair AI Regulation.

**To cite this article:** C. Kutterer, T. Karathanasis, The AI Act's GPAI Code: Hidden Policy Choices, AI Regulation Papers 25-03-1, [AI-Regulation.com](https://www.ai-regulation.com), March 7th, 2025.

These statements are attributable only to the author, and their publication here does not necessarily reflect the view of the other members of the AI-Regulation Chair or any partner organizations.

**This work has been partially supported by MIAI @ Grenoble Alpes, (ANR-19-P3IA-0003) and by the Interdisciplinary Project on Privacy (IPoP) of the Cybersecurity PEPR (ANR 22-PECY-0002 IPOP).**