Commentary and Critique **Research Commentary**
Vol. 10 (2026) · ISSN 3006-6220 journal.pe.ac/Commentary-and-Critique

Deferred Enforcement, Applicable Reporting: The General-Purpose AI Code of Practice in the EU AI Act's Enforcement Interval

Natalie Goh[a]

[a]Academy of Public Policy and Economics (APPE), London E1 0SG, United Kingdom

Corresponding author: Natalie Goh, natalie.goh@pe.ac.

Published 20 June, 2026.

DOI: https://doi.org/10.62432/APPE.CC.2026.10.26

Article webpage: https://journal.pe.ac/Commentary-and-Critique/article/view/26

Abstract

The EU AI Act's general-purpose AI provider obligations became applicable on 2 August 2025; Article 113 provides that the Commission's Chapter IX supervision and enforcement powers, and Article 101 on fines, apply from 2 August 2026. The General-Purpose AI Code of Practice was published into that interval. This article asks what adherence could do for a provider, and non-adherence do to one, during those 365 days, reading the answer off the statute and the signing arrangements rather than off signature counts. It identifies six consequences the Commission publicly attached to signing. Five operate wholly through Chapter IX, Section 5 or Article 101 and were unexercisable throughout; the sixth is partly an enforcement posture under Article 89(1), unexercisable for the same reason, and partly a reputational effect that no provision of the Regulation confers or withholds. Articles 56(5) and 56(6) sit in Chapter V and applied from 2 August 2025: participants report regularly to the AI Office on the implementation of their commitments, and the AI Office and the Board monitor and evaluate their achievement of the Code's objectives. The Code's Commitment 7 is grounded in Article 56(5) and sets a six-month Model Report cycle for the most capable models, which the Commission's adequacy Opinion assesses. Signature therefore allocated an applicable reporting and monitoring relationship that non-signature did not, rather than the enforcement-dependent benefits. Three adherence patterns were observable by the 1 May 2026 cutoff; the article establishes no provider's Article 111(3) transition status, no systemic-risk classification and no provider's reasons.

Keywords: EU AI Act, general-purpose AI models, code of practice, co-regulation, enforcement

1. Introduction

Article 113 creates an unusual interval from 2 August 2025 to 2 August 2026 for providers of general-purpose AI models in the European Union. The statutory obligations apply during that interval, while the authority charged with enforcing them has no applicable Chapter IX power to do anything about them. Into that interval the European Commission's AI Office placed the General-Purpose AI Code of Practice, a voluntary instrument published on 10 July 2025 and assessed as adequate on 1 August 2025, one day before the obligations took effect (European Artificial Intelligence Board, 2025; European Commission, 2025c, 2025e). That the AI Act would lean on co-regulatory instruments of this kind was identified before the Code appeared (Larouche, 2025).

© 2026 Academy of Public Policy and Economics · Open Access (CC BY 4.0).

Most commentary on the Code has concentrated on what its chapters require and on who signed. This article asks a narrower question with a checkable answer: given the statutory calendar, what could adherence to the Code actually do for a provider, or non-adherence do to one, during the interval in which the observed signature decisions were taken?

The answer separates two things that are easily run together. The consequences the Commission publicly advertised for signing and for not signing were almost all exercised through provisions that were not yet applicable, so the Commission could not confer or impose them during the interval. But the Code's own statutory basis, Article 56, sits in Chapter V and was advanced to 2 August 2025 along with the obligations; its reporting and monitoring paragraphs gave the AI Office an applicable relationship with participants throughout. Signature during the interval therefore did not buy the advertised enforcement benefits, and it did allocate a live reporting relationship. Each half is a claim about which articles of Regulation (EU) 2024/1689 were applicable when, and each follows from the text of Article 113 read against the Regulation's chapter structure. Nothing in either implies anything about regulatory laxity or about the Commission's intentions.

Section 2 sets out the statutory calendar, Section 3 maps what the Code covers, Section 4 traces each advertised consequence to its instrument, and Section 5 sets out the Article 56 channel that was applicable in the same period. Section 6 examines the signing instrument, Section 7 records the adherence patterns visible by the evidence cutoff of 1 May 2026, and Section 8 states what the split does and does not show.

2. The three dates set by Articles 111 and 113

Article 113 of the Regulation sets a general application date of 2 August 2026 and then carves out exceptions. The relevant exception reads: "Chapter III Section 4, Chapter V, Chapter VII and Chapter XII and Article 78 shall apply from 2 August 2025, with the exception of Article 101" (Regulation (EU) 2024/1689, art. 113, second para., point (b)). Chapter V contains the obligations for providers of general-purpose AI models, and it also contains Article 56, on codes of practice. Article 101 contains the Commission's power to fine those providers. So the obligations advanced by a year, and the fines did not.

The carve-out is wider than the express exclusion of Article 101 suggests. The list in Article 113 does not include Chapter IX, and Chapter IX, Section 5 is headed "Supervision, investigation, enforcement and monitoring in respect of providers of general-purpose AI models." It runs from Article 88 to Article 94 and contains Article 88, which gives the Commission exclusive powers to supervise and enforce Chapter V; Article 89, on monitoring; Article 91, the power to request documentation and information; Article 92, the power to conduct evaluations; and Article 93, the power to request measures up to withdrawal of a model from the market. Article 113 provides that none of these applies before 2 August 2026.

The Commission states the same conclusion in its own guidelines: "In the first year from 2 August 2025 onwards, the Commission cannot take any enforcement actions because its enforcement powers only enter into application on 2 August 2026. This lack of enforcement does not put into question the applicability of the obligations laid down in Chapter V of the AI Act from 2 August 2025 onwards" (European Commission, 2025b, para. 112). The Commission had said as much at publication, describing the rules as becoming "enforceable by the AI Office of the Commission one year later as regards new models and two years later as regards existing models" (European Commission, 2025e). Practitioner commentary has often narrowed that reading to a "grace period" for fines (Jones Day, 2025; Skadden, 2025).

Against that, one retained Commission document reads differently. The explanatory notice accompanying the training-content template, a Communication of 5 December 2025 and so written inside the interval, states at paragraph (26) that where the AI Office verifies a provider's summary "the AI Office has all enforcement powers under the AI Act", and that non-compliance "may be sanctioned with fines of up to 3% of the provider's annual total worldwide turnover in the preceding financial year or EUR 15 000 000,

whichever is higher" (European Commission, 2025d, para. 26). The sentence carries no temporal qualifier. Read against Article 113 it describes which powers the AI Office holds under the Regulation rather than the date from which it holds them, and the guidelines quoted above, which are the Commission's own considered statement on that date, are unambiguous in the other direction. The passage is recorded here because the calendar was not observed consistently even in the Commission's own drafting, which is a fact about the record this article reads and not a competing legal position.

A third date matters and is easy to lose. Article 111(3) provides that "[p]roviders of general-purpose AI models that have been placed on the market before 2 August 2025 shall take the necessary steps in order to comply with the obligations laid down in this Regulation by 2 August 2027." For a model already on the market when the obligations took effect, 2 August 2026 is when the Commission acquires its powers, and 2 August 2027 is when compliance is due. The provision sets a later compliance deadline; it does not defer the assessment. Figure 1 places the three bands on one axis, together with the dated Code milestones and the provider statements discussed in Section 7.

One further interval is worth recording. Article 56(9) required codes of practice to be "ready at the latest by 2 May 2025." The Code appeared on 10 July 2025, 69 days later, and was assessed as adequate 22 days after that, one day before Chapter V became applicable. The statute's own fallback is dated differently: the second subparagraph of Article 56(9) allows the Commission to provide common rules by implementing act if a code "cannot be finalised" by 2 August 2025, which the Code preceded by 23 days. The instrument arrived after the preparation date in the first subparagraph and before the operative trigger in the second, and was confirmed on the eve of the obligation it was meant to help satisfy.

3. What the Code covers

The Code consists of three separately authored chapters (Oliver & Bommasani, 2025; Peukert & Castets-Renard, 2025; Samwald et al., 2025). The Commission counts twelve commitments across them: one in Transparency, one in Copyright, and ten in Safety and Security (European Commission, 2025c, para. 7). Recounting the measure identifiers directly from the three published chapters gives three measures under the Transparency commitment, five under Copyright, and thirty-two across the Safety and Security commitments, forty in total. The Transparency chapter states that the Annex XI, Section 2 documentation items for systemic-risk models are handled by Measure 10.1 of the Safety and Security chapter, so that measure is counted once, in the chapter that contains it.

Table 1 maps the statutory obligations onto that structure.

Two qualifications belong with the table. The first concerns who the first two obligations reach. Article 53(2) provides that the obligations in Article 53(1), points (a) and (b), "shall not apply to providers of AI models that are released under a free and open-source licence that allows for the access, usage, modification, and distribution of the model, and whose parameters, including the weights, the information on the model architecture, and the information on model usage, are made publicly available," and that the exception does not apply to models with systemic risk. The Transparency chapter carries the same limit in its own scope note and in its Commitment 1 legal-text box, and the Commission repeats it in its Q&A (European Commission, 2025f). The exception matters for what refusal signifies: for a provider whose models meet those conditions and do not present systemic risk, two of the four Article 53(1) obligations do not apply at all, so the Code has correspondingly less to demonstrate. This study did not determine whether any provider in Table 2 meets the Article 53(2) conditions for any model, and makes no claim that any does.

The second concerns the one obligation the table records as uncovered. Article 56(2)(b) had listed the level of detail for the training-content summary as a matter for a code of practice, but the AI Office prepared the template separately, and the Copyright chapter refers to it rather than specifying it. The

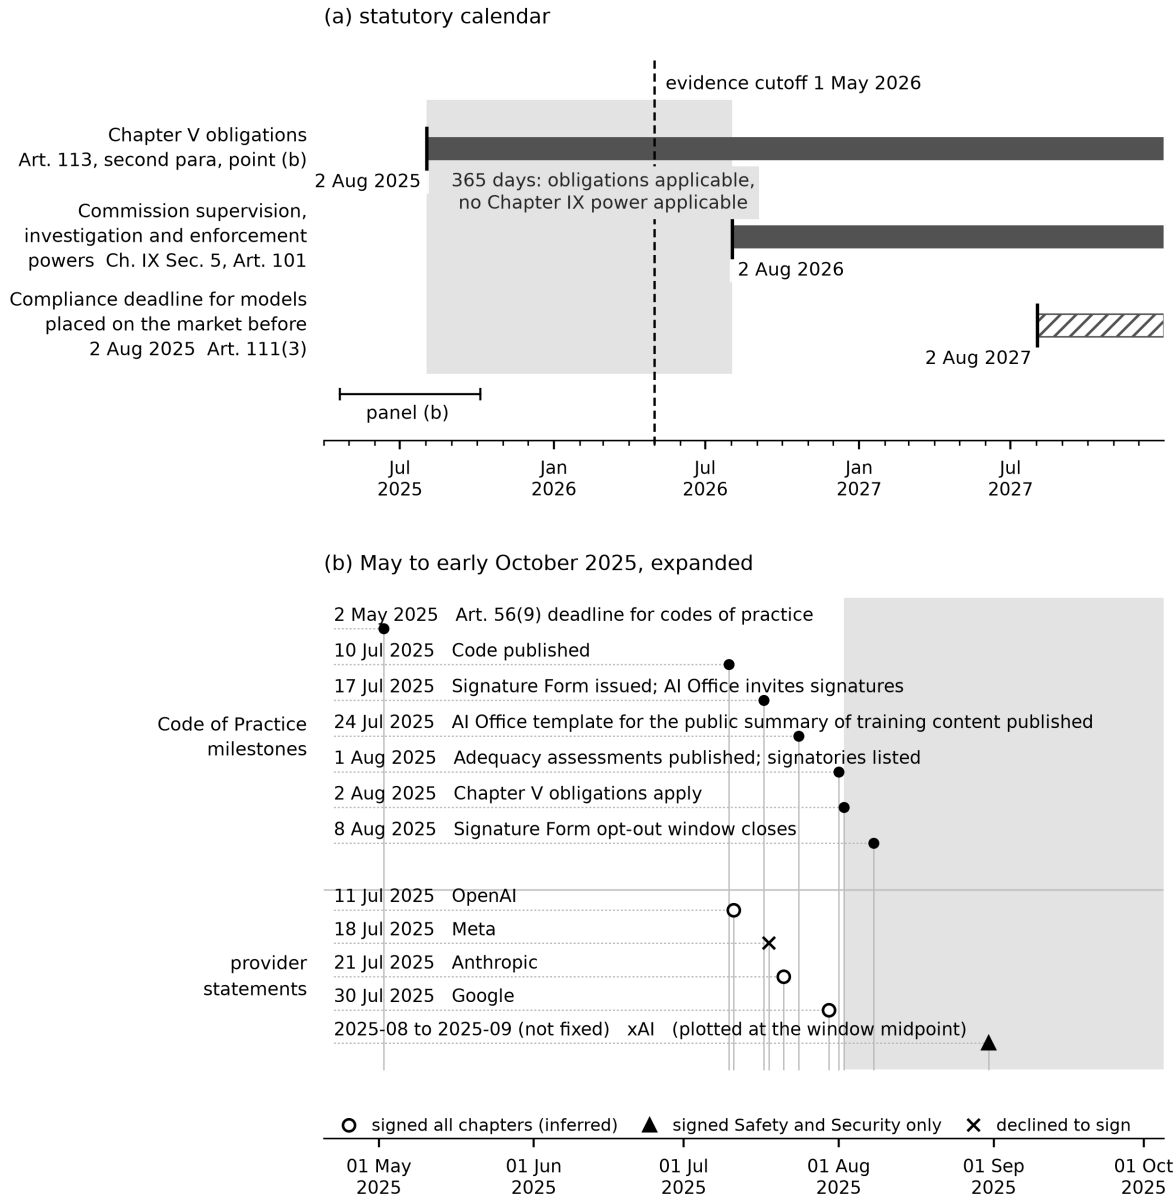


Figure 1. The statutory calendar for providers of general-purpose AI models. Panel (a): the three bands set by Article 113, second paragraph, point (b), by Article 111(3) and by the general application date, with the 365-day interval shaded. The two applicability bands are drawn solid and the Article 111(3) band is hatched, because it opens on a date by which compliance is due rather than on a date from which provisions apply. Panel (b): May to early October 2025 expanded, because every dated Code of Practice milestone and every provider statement examined in this article falls inside it; the panel runs five days past the end of the coded August–September window, and its title is set from that extent. The shading marks the days on which the Chapter V obligations were already applicable. The xAI marker is plotted at the midpoint of the coded August–September 2025 window, computed by the script from that window; the evidence does not fix the date (see Section 7). Milestones are read from analysis/data/timeline-v03.csv and provider rows from analysis/data/adherence-observations-v03.csv; plotted dates are listed in figures/figure1-calendar-data-v03.csv; vector source: figures/figure1-calendar-v03.pdf. The Chinese-language render is figures/figure1-calendar-v03.zh.png; both are produced by the same script and the same data and differ only in the printed language.

Table 1. Coverage of the Article 53 and Article 55 obligations by the three chapters of the Code of Practice. The Safety and Security measures are listed as a shared block because the chapter's thirty-two measures serve the four Article 55(1) obligations jointly rather than being divided among them. The addressee column records the Article 53(2) exception; the "Covered" column records whether any Code measure addresses the obligation, which is a narrower question than the Commission's adequacy conclusion discussed below. Generated by analysis/build_tables_v03.py from analysis/data/obligation-coverage-v03.csv.

AI Act provision	Obligation	Addressees	Code chapter	Measures	Covered
Art. 53(1)(a)	Technical documentation for the AI Office and national competent authorities (Annex XI, Sec. 1)	all providers, except free and open-source models under Art. 53(2)	Transparency	1.1–1.3	yes
Art. 53(1)(b)	Information and documentation for downstream providers (Annex XII)	all providers, except free and open-source models under Art. 53(2)	Transparency	1.1–1.3	yes
Art. 53(1)(c)	Policy to comply with Union copyright law, including Art. 4(3) DSM reservations	all providers	Copyright	1.1–1.5	yes
Art. 53(1)(d)	Public summary of the content used for training	all providers	none	—	**no**
Art. 55(1)(a)	Model evaluation, including documented adversarial testing	systemic-risk providers	Safety and Security	1.1–10.2 (32, shared)	yes
Art. 55(1)(b)	Assessment and mitigation of systemic risks at Union level	systemic-risk providers	Safety and Security	1.1–10.2 (32, shared)	yes
Art. 55(1)(c)	Tracking, documenting and reporting serious incidents	systemic-risk providers	Safety and Security	1.1–10.2 (32, shared)	yes
Art. 55(1)(d)	Cybersecurity protection for the model and its infrastructure	systemic-risk providers	Safety and Security	1.1–10.2 (32, shared)	yes
Annex XI, Sec. 2	Additional documentation for systemic-risk models	systemic-risk providers	Safety and Security	10.1	yes

Commission records the decision plainly: the Code "will refer to the AI Office's template for the summary that will specify the adequate level of detail, as required in Article 53(1), point (d), of the AI Act, regardless of whether the providers adhere to the Code" (European Commission, 2025c, para. 27). The template exists and is dated. The Commission's page carries a publication date of 24 July 2025, and the instrument published on that date is Communication C(2025) 5235 final, identified by number in Peukert's (2026) note 88; the original was not obtained for this study. The version retained here is the successor Communication C(2025) 8311 final of 5 December 2025, whose explanatory notice states that the Article 53(2) open-source exception "does not apply to the obligation to make publicly available the Summary" and that providers of models placed on the market before 2 August 2025 should publish the corresponding summary no later than 2 August 2027 (European Commission, 2025d). No claim is made here about differences between the two versions. Peukert (2026) treats the in-house preparation of the template as evidence of the AI Office's controlling role in a nominally expert-led process.

The Commission's own conclusion runs the other way at the level of the obligation, and it should be stated rather than worked around. Four paragraphs after the sentence quoted above, the Opinion assesses that the Transparency and Copyright chapters "adequately cover the obligations provided for in Article 53(1), and the parts thereof outlined in Article 56(2), points (a) and (b)," and the concluding paragraph finds that the Code "adequately covers the obligations provided for in Articles 53 and 55 of the AI Act and

meets the aims according to Article 56 of the AI Act" (European Commission, 2025c, paras. 28, 57). The two statements are compatible on a narrow reading, which this article adopts: the Commission is assessing whether the Code adequately covers the Article 53(1) obligations taken together, given that the level of detail for the summary is fixed elsewhere by an instrument that binds signatories and non-signatories alike, while Table 1 records the narrower fact that no Code measure addresses Article 53(1)(d). On that reading a provider that signs every chapter has still not addressed one of its four Article 53(1) obligations through the Code, but it has not thereby fallen short of what the Commission assessed the Code to cover.

4. Six advertised consequences and the date they share

The Commission has been specific about what signing does. For signatories, it "will focus its enforcement on monitoring their adherence to the code, which offers greater predictability and reduced administrative burden" (European Commission, 2025a); signatories "will therefore benefit from increased trust from the Commission and other stakeholders" (European Commission, 2025b, para. 91); and commitments implemented in line with the Code may be taken into account "as a mitigating factor when fixing the amount of fines, depending on the specific circumstances (Article 101(1) AI Act)" (European Commission, 2025b, para. 93). Providers that do not adhere "must demonstrate compliance via other adequate means and will have to report the measures they have implemented to the AI Office," are expected to explain their adequacy "for instance by carrying out a gap analysis that compares the measures they have implemented with the measures set out by the code," and "may also be subject to a larger number of requests for information and requests for access to conduct model evaluations throughout the entire model lifecycle" (European Commission, 2025g).

Six is this article's parse across the Commission's pages, not a Commission enumeration. Four further advertised items are recorded in analysis/data/pathway-consequences-v04.csv and deliberately not counted, each with its reason. The page that carries the first quotation above says that providers who sign "will enjoy streamlined compliance with the obligations for general-purpose AI models in the AI Act," which names the demonstration route itself rather than a further consequence, and that "[s]ignatories will be publicly listed on 1 August 2025" (European Commission, 2025a). The press release of 10 July 2025 carries a fourth variant of the signatory benefits, promising "a reduced administrative burden and increased legal certainty compared to providers that prove compliance in other ways" (European Commission, 2025e); the administrative-burden element is already counted below, and legal certainty relative to other demonstration routes restates the predictability of the enforcement posture rather than naming a further power. A Commission spokesperson told a trade outlet that companies complying by other means "may be exposed to more regulatory scrutiny by the AI Office," which glosses the information-request and evaluation-access item listed above, in a press report rather than on a Commission page (Claburn, 2025). A different parse of the same pages would produce a different count; the file records each item, its source and its verbatim basis so that the parse can be checked.

Five of the six run through a specific deferred provision. Focused enforcement and the monitoring of adherence are Article 89(1), which empowers the AI Office to monitor compliance "including their adherence to approved codes of practice." For providers that do not adhere, the Commission's sentence runs two duties together and they part company on the calendar: the duty to demonstrate compliance by other adequate means is imposed by the third sentences of Articles 53(4) and 55(2), which are in Chapter V and applied from 2 August 2025, while reporting the implemented measures to the AI Office and producing the expected gap analysis are the practical form of the supervisory relationship under Articles 88 and 89. Only the reporting half is counted among the six, and the data file isolates it. Heightened exposure to information requests and evaluation access is Articles 91 and 92. Fine mitigation is Article 101(1). Articles

88 to 94 are Chapter IX, Section 5; Article 101 is expressly excluded from advance application. Article 113 provides that all five apply from 2 August 2026.

The sixth is of mixed character, and counting it whole would overstate the result. ”Reduced administrative burden” describes the practical effect of the enforcement posture in the first item, so it shares that item’s date. ”Increased trust from the Commission and other stakeholders” does not divide so cleanly: trust from the Commission is an attitude within the supervisory relationship, and trust from other stakeholders is a reputational effect that no provision of the Regulation confers or withholds and that Section 8 places outside this article’s evidence. So the accurate statement is not a clean six for six. Five of the six advertised consequences were wholly unexercisable throughout the interval, because the provisions that carry them were not applicable; the sixth was unexercisable in its statutory part for the same reason, and in its reputational part it lies outside the statutory analysis altogether. The coded data carries the same split: the five deferred items each hold one application date, and the mixed item holds a date only for its statutory part. Figure 2 sets out the two pathways, the six consequences in parallel, and what the interval did and did not leave exercisable.

The coupling of the deferred powers to one another appears inside the drafting of the powers themselves. Article 91(4) requires that a request for information ”state the legal basis and the purpose of the request, specify what information is required, set a period within which the information is to be provided, and indicate the fines provided for in Article 101 for supplying incorrect, incomplete or misleading information.” The information-gathering power is drafted so that its formal notice points at a penalty provision, and during the interval that penalty provision was itself inapplicable.

One term in the chain already has a literature, and it reaches a conclusion. Article 89(1) empowers monitoring of compliance ”including their adherence to approved codes of practice,” and the third sentences of Articles 53(4) and 55(2) impose the alternative-means duty on providers ”who do not adhere to an approved code of practice.” Peukert (2026) gives the term a titled subsection and observes that ”it is doubtful who is supposed to bring about this approval and what the procedure should be,” because Article 56(6) ”provides for two different procedures”: an adequacy assessment published by the AI Office and the Board under the first subparagraph, and approval by implementing act giving the code ”a general validity within the Union” under the second. The wording, he notes, ”points to Article 56(6) subpara. 2,” which uses the verb ”approve,” and the Commission’s Q&A appears to him to share that view; the Commission’s guidelines say instead that providers demonstrate compliance by adhering to a code of practice ”that is assessed as adequate by the AI Office and the Board (Article 56(6) AI Act)” (European Commission, 2025b, para. 91). He resolves the term the other way. Reading ”approved” as ”assessed as adequate” ”brings the AI Act in line with the regulation of CoCs under the DSA and the GDPR,” which distinguish an opinion on a code’s appropriateness from a separate Commission decision granting an approved code general validity, and it mirrors the two-step structure of Article 56(6) itself, in which limited legal effect attaches to the adequacy assessment and general validity to the implementing act. His conclusion is stated directly: ”There are thus good reasons to assume that a CoP acquires legal effect and thereby becomes part of EU law through an adequacy assessment according to Article 56(6) subpara. 1 AI Act.” The provisions he collects as turning on approval — Articles 53(4) third sentence, 55(2) third sentence and 89(1) — are the same three assembled here.

What happened on 1 August 2025 was an assessment of adequacy under the first subparagraph, published by the Commission and by the Board (European Artificial Intelligence Board, 2025; European Commission, 2025c); nothing located in this study shows that the implementing act was adopted. This article adopts neither reading. What it adds is that the calendar result holds under both. On Peukert’s reading the Code acquired legal effect on 1 August 2025 and the words ”approved code of practice” were satisfied throughout the interval — and the consequences the Commission advertised for satisfying them still sat in Articles 88 to 94 and in Article 101, which did not apply until 2 August 2026. On the stricter reading

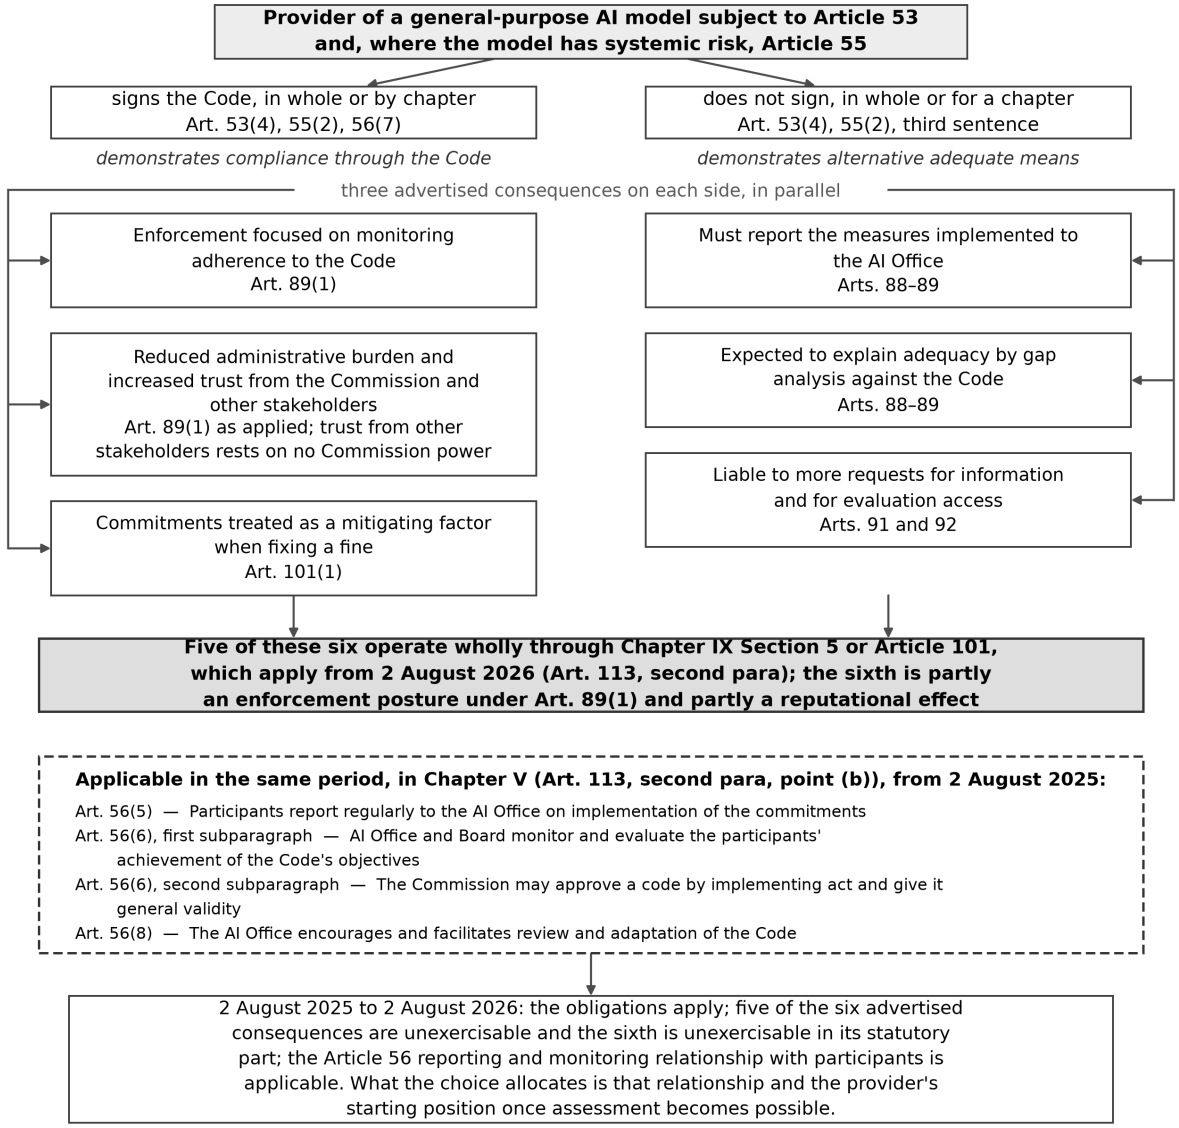


Figure 2. The two compliance pathways under Articles 53(4) and 55(2), the advertised consequences attached to each, and the applicable Article 56 channel. Within each column the three consequences are independent of one another and are drawn from a shared stem for that reason. The summary boxes are written from the coded rows rather than typed, so the figure states the same five-plus-one split as the text: five consequences wholly unexercisable during the interval, the sixth unexercisable in its statutory part. Consequences, instruments, application dates, source identifiers and verbatim source text are taken from analysis/data/pathway-consequences-v04.csv; the parse is summarised in results/pathway-consequences-check-v04.md; vector source: figures/figure2-mechanism-v03.pdf. The Chinese-language render is figures/figure2-mechanism-v03.zh.png; both are produced by the same script and the same data and differ only in the printed language.

the words were not satisfied at all before an implementing act, and the advertised consequences stood further from the signatory still. That is why the claims made here are stated in terms of the Commission's advertised consequences rather than of a settled reading of "approved."

What remains in the interval is worth stating precisely, because it is not nothing. The obligations in Articles 53 and 55 applied in full from 2 August 2025; the Commission says so in the same passage in which it says it can take no enforcement action (European Commission, 2025b, para. 112). A provider that ignored them was in breach of an applicable obligation, and private litigation is not subject to the Commission's calendar (Jones Day, 2025). Section 5 sets out the further point that the Code's own statutory basis was applicable throughout.

This matters for how the Code's legal effect is described. Articles 53(4) and 55(2) make adherence a way to "demonstrate compliance" and reserve the presumption of conformity for harmonised standards; the Commission's guidance restates the contrast (European Commission, 2025b, para. 97), its Q&A says adherence is a means of demonstrating compliance "while not providing a presumption of conformity with the AI Act" (European Commission, 2025f), and the Code's own Objectives state that adherence "does not constitute conclusive evidence of compliance." Peukert (2026) draws the same distinction. The point still needs making, because it is not uniformly observed: an analysis published by the Center for Security and Emerging Technology on 30 July 2025 described the Code as offering a "presumption of conformity" and treated it as "a strong incentive for model providers to sign," glossing it as regulators assuming that adherence demonstrates compliance (Hoffmann, 2025). Combined with the timing, the overstatement compounds: an instrument that confers a weaker status than a harmonised standard also conferred none of its advertised enforcement-linked consequences during the first year of its operation.

5. The channel that was applicable: Articles 56(5) to 56(8)

Article 56 is the Code's statutory basis, and it is in Chapter V. Article 113, second paragraph, point (b) advances the whole of Chapter V to 2 August 2025 and excludes only Article 101, so Article 56 applied in full throughout the interval, all nine of its paragraphs. Four of them bear on the question asked here, and two of those put the AI Office and the signatories into a stated relationship with one another.

Article 56(5) provides that "[t]he AI Office shall aim to ensure that participants to the codes of practice report regularly to the AI Office on the implementation of the commitments and the measures taken and their outcomes, including as measured against the key performance indicators as appropriate," and that reporting commitments "shall reflect differences in size and capacity between various participants." Article 56(6), first subparagraph, provides that "[t]he 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 second subparagraph carries the approval power discussed in Section 4, and Article 56(8) provides that the AI Office "shall, as appropriate, also encourage and facilitate the review and adaptation of the codes of practice, in particular in light of emerging standards."

None of this stayed at the level of the statute. The Code's Safety and Security chapter grounds its reporting commitment in Article 56(5) by name: the legal-text box above Commitment 7 reads "LEGAL TEXT: Articles 55(1) and 56(5) AI Act," and the commitment binds signatories to produce a Safety and Security Model Report before placing a model on the market, to keep it up to date and to notify the AI Office of it. Measure 7.6 adds that, for a model "amongst their respective most capable models available on the market," signatories "will provide the AI Office with an updated Model Report at least every six months," subject to three stated exceptions. The Commission's adequacy Opinion assesses exactly this material under the heading that names Article 56(4) and (5), recording that "participants commit to reporting to the AI Office information about their model and adherence to the Safety and Security Chapter pursuant to Commitment 7," that Measures 1.1 and 1.4 carry a further list of systemic-risk information to

be compiled and reported, and that proactive reporting is not required for the Transparency and Copyright chapters because the obligations those chapters cover already produce documentation (European Commission, 2025c, paras. 42, 43). The pre-cutoff Q&A cites Article 56(5) for the proposition that the Code's reporting commitments "reflect differences in size and capacity between various providers" (European Commission, 2025f).

What that meant inside the interval can be stated concretely. A provider that signed the Safety and Security chapter in July or August 2025 entered a reporting relationship with the AI Office. That relationship had an applicable statutory basis, it was written into the Code text the provider signed, the Commission assessed it in the act that made the Code operative, and for the provider's most capable models it ran on a stated six-month cycle.

The same question put on the other side does not come back empty. A provider that did not sign had no Article 56(5) channel, because Article 56(5) addresses participants in a code of practice. It carried applicable Chapter V duties all the same. The third sentences of Articles 53(4) and 55(2) provide that providers "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," and Article 53(3) requires every provider to "cooperate as necessary with the Commission and the national competent authorities in the exercise of their competences and powers pursuant to this Regulation." Both are in Chapter V, and both applied from 2 August 2025. What was deferred on that side was the machinery rather than the duty: the Commission's assessment of the means demonstrated, and the information-gathering that would put those means in front of it, are exercised under Articles 88, 89, 91 and 92, which is why the reporting expectation the Commission attached to the alternative-means route was not exercisable inside the interval. The asymmetry the interval created is therefore not consequence against no consequence but a difference in kind on each side. Signature allocated an applicable reporting and monitoring relationship and did not buy the enforcement-dependent benefits; refusal avoided that relationship, left the Chapter V duty to demonstrate alternative adequate means in place, and did not attract the advertised additional burdens.

Two further things follow from the same paragraphs. The Commission's reserved power to revisit its own assessment, which Section 6 relies on, is itself exercised under applicable provisions: the Opinion states that it does not "pre-empt the Commission from assessing the Code as inadequate in the future following its regular monitoring and evaluation of the achievement of the Code's objectives and its contribution to the proper application of the AI Act pursuant to Article 56(6) of the AI Act or following the Commission's review and adaptation of the Code pursuant to Article 56(8) of the AI Act" (European Commission, 2025c, para. 3). And one advertised item is attached to a date inside the interval. The invitation page stated that "[s]ignatories will be publicly listed on 1 August 2025" (European Commission, 2025a). Public listing is a Commission act attached to the signing decision, it invokes no supervisory power, and the date it names is the day before the obligations took effect. This study did not obtain a copy of the list as published on that date, so what the record supports is the Commission's stated undertaking rather than independent confirmation that the listing occurred then. The item is not among the six because it is not an enforcement-linked consequence, but it is the one advertised effect of signing that the Commission tied to a date it could reach without any Chapter IX power.

6. A commitment fixed to one text

The signing instrument is a single page, and its three operative conditions shape what a signature meant. The form states that it "will only become binding once the GPAI Code of Practice is assessed as adequate by the European AI Office and the European Artificial Intelligence Board pursuant to Article 56(6) AI Act, and at the earliest with the entry into application of Chapter V of the AI Act on 2 August

2025"; that the signatory "reserves the right to opt out of signing the GPAI Code of Practice at any point before or within 7 days of the publication of the adequacy assessments under Article 56(6) AI Act by emailing the European AI Office"; and that it "only binds as to the text of the GPAI Code of Practice published on July 10, 2025," with any later updates to be "shared with the Signatories for their consideration" (European AI Office, 2025).

The adequacy assessments were published on 1 August 2025 and Chapter V applied from 2 August 2025, so the seven-day reservation ran until 8 August 2025. For six days after the obligations took effect, a signature had not yet passed the last point at which the signatory could decline to be bound at all. That is a narrower statement than it may appear, and the narrowing matters. The Form's clause is a right to opt out of signing, which prevents the Signature Form from becoming binding; it is not the only route out. The Commission's signing Q&A answers the question "Can a signatory withdraw their signature?" with "Yes, a signatory may withdraw their signature to the Code at any time" (European Commission, 2025g). Revocability did not end on 8 August 2025. What ended on that date was the window in which a signatory could still prevent the commitment from ever taking effect, at a moment when the obligations the commitment was meant to help demonstrate compliance with had already been applicable for six days. OpenAI's announcement was drafted to the same conditionality, stating an intention to sign "subject to the current version being formally approved by the AI Board during its upcoming adequacy assessment" (OpenAI, 2025).

The version clause has a longer reach. The Commission intends to "consider facilitating formal updates to the Code at least every two years, for instance based on the emergence of standards, relevant technological developments, or changes in the risk landscape," and records that its adequacy opinion does not bind it in interpreting the AI Act and does not pre-empt a later finding of inadequacy (European Commission, 2025c, paras. 3, 51). A signatory is bound to a fixed text and must agree afresh to any revision; the Commission may revise its assessment of that text without the signatory's agreement, under Articles 56(6) and 56(8). The instrument is not symmetrical, and the asymmetry runs in the direction that makes signature cheaper.

7. What was observable by the cutoff

Table 2 records the adherence patterns for which dated evidence was obtained before the 1 May 2026 cutoff.

The xAI entry rests on a Commission statement recorded before the cutoff rather than on the Commission's live page. Peukert (2026), published online on 26 March 2026, quotes the Commission as announcing that "xAI signed up to the Safety and Security Chapter; this means that it will have to demonstrate compliance with the AI Act's obligations concerning transparency and copyright via alternative adequate means." The page identification in that article's notes is not straightforward and is recorded here as it stands: the note attached to the quotation cross-refers to the Commission's press-release page of 10 July 2025, last accessed 16 October 2025, while the Commission page that carries the sentence in the copy consulted for this study is the continuously updated Code of Practice policy page, which the same article cites elsewhere in its 9 September 2025 version, last accessed 27 September 2025. The copy of the press-release page retained here, whose last-update date is 24 September 2025, does not contain the sentence. What the evidence supports is therefore that the Commission had made the statement in a public page by October 2025, as recorded before the cutoff by a peer-reviewed source that quotes it verbatim; which page carried it is not settled here. The live policy page consulted for this study carries a last-update date after the cutoff and is not relied on for any finding. The exact date of the xAI signature is not established: the Commission's wording is undated, and this study located no dated primary statement by xAI.

Table 2. Observed adherence patterns and their evidential basis. The chapter cells for the three full-signature providers record announced intentions to sign together with an inferred chapter scope, which is the strength the claim ledger permits for those sources: each is the provider's own public position on the stated date, not evidence of an executed signature, and the chapter scope is an inference from the absence of a chapter qualifier in the Commission wording quoted in Peukert's 2026 analysis. Rows are in chronological order; the coded window sorts last. Generated by analysis/build_tables_v03.py from analysis/data/adherence-observations-v03.csv.

Provider	Pattern	Transparency	Copyright	Safety and Security	Statement date	Evidential basis
OpenAI	full	announced (inferred)	announced (inferred)	announced (inferred)	2025-07-11	own post announcing an intention to sign; chapter scope inferred from the absence of a chapter qualifier in the Commission wording quoted in Peukert's 2026 analysis
Meta	refusal	—	—	—	2025-07-18	own post, quoted verbatim in two same-day reports
Anthropic	full	announced (inferred)	announced (inferred)	announced (inferred)	2025-07-21	own post announcing an intention to sign; chapter scope inferred as for OpenAI
Google	full	announced (inferred)	announced (inferred)	announced (inferred)	2025-07-30	own post announcing it will sign; chapter scope inferred as for OpenAI
xAI	partial	—	—	signed (Commission statement)	2025-08 to 2025-09 (not fixed)	Commission statement quoted in Peukert (2026); date not independently fixed

The chapter scope for the three full signatories is an inference, and a bounded one: the Commission attaches a chapter qualifier to xAI and, in the wording Peukert reports, none to the others. The three providers' own posts announce intentions to sign and do not name chapters. The one official dated list obtained is the January 2026 Vademecum of the Signatory Taskforce, which names twenty-three participating signatories and does not include xAI (European Commission, 2026). That document states that the Taskforce "is comprised of the attached list of Signatories as 'Members'" and that "Signatories that are not Members may communicate at any time their intent to join the Taskforce," so twenty-three is a lower bound on the number of signatories at that date rather than a count of them. The same secondary carrier that supplies the xAI sentence supplies a higher figure, recorded here on the same terms: Peukert (2026) reports that "[t]wenty-six companies of different sizes responded to this invitation," naming Anthropic, Google, Microsoft and OpenAI among them together with Aleph Alpha and Mistral, and cites the Code of Practice policy page in its 9 September 2025 version for the figure. That is a pre-cutoff secondary report of what a continuously updated page said on a stated date, not an official list obtained for a stated date, and it is used here only as such.

The stated reasons vary in a way the legal analysis should not flatten. Anthropic framed signature as continuous with its existing practice, noting that the Code's Safety and Security Frameworks "build upon Anthropic's own Responsible Scaling Policy" (Anthropic, 2025). Google signed while objecting to the

substance, warning that ”departures from EU copyright law, steps that slow approvals, or requirements that expose trade secrets could chill European model development and deployment” (Walker, 2025). Meta declined on grounds of legal uncertainty: ”We have carefully reviewed the European Commission’s Code of Practice for general-purpose AI (GPAI) models and Meta won’t be signing it. This Code introduces a number of legal uncertainties for model developers, as well as measures which go far beyond the scope of the AI Act” (Kaplan, 2025, as cited in Claburn, 2025, and in Subin, 2025). Google’s position shows that full signature and substantive objection to the text were compatible in practice. These are the reasons the providers gave publicly; this article does not treat them as a complete account of why any provider decided as it did.

That chapter-selective signature was available is established independently of the xAI instance. The AI Office’s signing Q&A, dated 17 July 2025, states that ”[a]ny opt-out from chapters of the code of practice results in losing the benefits of facilitating the demonstration of compliance in that respect,” and separately that providers without systemic risk ”may optionally sign the Safety and Security chapter” and are then expected to specify the models covered (European Commission, 2025g). Peukert (2026) works through the resulting legal question and concludes that Article 56(7), which allows providers of models without systemic risk to limit adherence to the Article 53 chapters, governs the inverse of the xAI case, so that other selective signatures should also be permissible and the Code operates as three de facto separate codes. That account is a participant’s: the author discloses having co-chaired the working group responsible for the copyright rules from September 2024 to July 2025.

8. What the split shows, and what it does not

Three patterns appeared, and each corresponds to a distinct position on both halves of the answer. Full signature produces a documented set of commitments against which Article 89(1) provides for adherence to be monitored from 2 August 2026, and, for the Safety and Security chapter, an applicable reporting relationship with the AI Office under Article 56(5) from 2 August 2025. Signature of the Safety and Security chapter alone produces both of those for the Article 55 obligations and leaves the Article 53(1)(a) to (c) obligations to be demonstrated by alternative adequate means. Refusal leaves all of them in that second category, with the reporting and gap-analysis expectations the Commission has described, none of which was exercisable during the interval. In every case the Article 53(1)(d) summary stands outside the Code and follows the AI Office template.

Article 111(3) then shifts the horizon again, though only for the models it reaches, and the reading of its scope is contested in a way that should be set out rather than resolved by assertion. On the text, the provision attaches to ”general-purpose AI models that have been placed on the market before 2 August 2025,” which is a statement about individual models and their placing-on-market dates rather than about providers or model series. The Commission reads the same words more broadly. Its guidelines state that it ”considers the lifecycle of a general-purpose AI model to begin at the start of the large pre-training run,” that ”[a]ny subsequent development of the model downstream of this large pre-training run performed by the provider or on behalf of the provider, whether before or after the model has been placed on the market, forms part of the same model’s lifecycle rather than giving rise to new models,” and that ”[a] model is thus considered to be the same model along its entire lifecycle” (European Commission, 2025b, para. 22). It applies that doctrine to the transition directly: after quoting Article 111(3) it adds that the provision ”covers such models throughout their entire lifecycle” (European Commission, 2025b, para. 106). On the Commission’s reading, a later version derived from the same pre-training run is the same model and is within the transition; on the narrower textual reading, a later release is outside it unless it independently qualifies. This article states both and adopts neither for any named provider, because the question that matters here does not turn on it: whichever reading applies, the transition is defined

by individual models and their pre-training runs, and it does not extend to a provider as a whole or to a model series by virtue of the provider's identity. This study established no placing-on-market date and no pre-training-run date for any particular model, and makes no claim about which of the models in Table 2 fall within the transition. The structural point stands without that determination: wherever the transition applies, the distance between a signature made in July 2025 and the moment when the deferred consequences bear on the model is two years rather than one, and a signatory may face both horizons at once across a portfolio. Nothing in the sources read here applies the Article 111(3) distinction to the observed signature decisions; Peukert's (2026) account does not mention Article 111 or 2027.

Several things the split does not show should be stated. It does not show that signatories are complying and non-signatories are not: adherence is a means of demonstrating compliance, not compliance, and the Code says so. It does not show that refusal is unlawful, since Articles 53(4) and 55(2) expressly preserve alternative adequate means. It does not show that the division tracks systemic-risk classification. This study located no provider disclosure of cumulative training compute against the 10^25 floating-point-operation threshold in Article 51(2), and no published list under Article 52(6) for any date before the cutoff; neither absence was searched exhaustively, and neither is asserted as a fact about the world. One retained source carries a third party's estimate rather than a provider disclosure: the same trade report that quotes the Meta statement in Table 2 describes Meta's Llama 4 Behemoth as trained at 5e25 floating-point operations (Claburn, 2025). That figure is the outlet's, not the provider's, and it is not used here to place any model above or below the threshold. That threshold is in any case a rebuttable presumption subject to the notification and contestation procedure in Article 52, and notifications under Article 52(1) are made to the Commission rather than published. Which providers are legally within the Article 55 regime is therefore not established by anything in this article, and it cannot be inferred from prominence or model size. Five cases are too few to support a claim about the population of providers, and the sample is drawn from providers whose statements were publicly reported, which favours large firms.

The interval was not free of consequence, and the evidence does not support saying that it was. The obligations themselves applied throughout; 272 of the 365 days had elapsed at this study's cutoff, with 93 remaining. The Article 56 reporting and monitoring relationship was applicable and was described in the instrument signatories signed, and the Chapter V duty to demonstrate alternative adequate means was applicable to those who did not sign. National and private-law exposure ran on its own timetable. Commercial and reputational effects are outside this article's evidence entirely, and their absence from the analysis is a limit on the analysis rather than a finding about them. What the evidence does support is that the enforcement-dependent inducements the Commission advertised for signing, and the particular additional burdens it advertised for not signing, described a supervisory regime that was not yet exercisable when providers acted on them. Applicability is also only the first condition: the harder powers, such as restricting a model's availability on the Union market, have been described as difficult to exercise outside clear non-compliance (CDT Europe, 2026).

9. Conclusion

The question was what the Code could constrain during the interval between applicable obligations and applicable enforcement. The answer is about rights and dates, and it has two parts.

Five of the six consequences this article identifies in the Commission's own pages are conferred by Chapter IX, Section 5 or by Article 101, and Article 113 provides that those provisions apply only from 2 August 2026; the sixth is partly an enforcement posture under the same deferred provisions and partly a reputational effect that no provision of the Regulation confers. Throughout the window in which the observed signature decisions were taken and published, the Commission therefore held no applicable power to grant the enforcement-dependent benefits or to impose the advertised additional burdens. That

is a statement about the Commission's competence on given dates, demonstrable from Article 113 read against the Regulation's chapter structure, with the Commission's own guidelines agreeing.

At the same time, Article 56 was applicable. Participants owed regular reports to the AI Office under Article 56(5) from 2 August 2025, the AI Office and the Board were to monitor and evaluate their achievement of the Code's objectives under Article 56(6) from the same date, and the Code's Commitment 7 and the Commission's adequacy Opinion put concrete content into both. Signature during the interval therefore allocated an applicable reporting and monitoring relationship, on a stated six-month cycle for the most capable models, that non-signature did not. Non-signature left in place a different applicable Chapter V duty, to demonstrate alternative adequate means, whose assessing end was not yet open. What signature created was fixed to the 10 July 2025 text, and it rested on an assessment the Commission may revisit under Articles 56(6) and 56(8).

The practical consequence is for how the early record should be used. Signature counts from the first year cannot measure the disciplinary effect of deferred enforcement powers, because throughout the period they cover the Commission held no applicable power to exercise the enforcement-linked consequences it had attached to signing. They can bear on a different and more tractable question: whether the Article 56 reporting relationship, which was applicable from the first day, produced documentation and risk practice that the alternative-means route did not. The Model Reports owed under Commitment 7 on a six-month cycle, and the AI Office's monitoring and evaluation under Article 56(6), are where an answer to that question would come from. The 2026 and 2027 records can address what the enforcement powers did once they applied; the 2025 record cannot, and should not be cited as though it could.

10. Data and materials

The statutory text was retrieved from the EU Publications Office resource interface for CELEX 32024R1689, which returns the Official Journal XHTML (L series, 2024/1689, 12.7.2024); the EUR-Lex front end returned an automated-access challenge for every entry point attempted. All Commission documents, Code chapters, the Signature Form, the training-content template and the Vademecum were downloaded from the Commission's own servers. The Peukert (2026) article is retained in the materials as an open-access PDF, so that the claims resting on it can be checked against the text rather than against a description of it.

Continuously updated Commission pages are not used to describe the position before the 1 May 2026 cutoff; where such a page carries a statement relied on here, the statement is taken from a dated pre-cutoff source that quotes it. The source register in sources/00-SOURCE-REGISTER-v03.md records each item's publication or version date, its position relative to the cutoff, the extent of text actually read, and the specific claim it supports; sources/SOURCE-HASHES-v03.json records a SHA-256 digest for each of the sixty-three retained source files. The claim-level ledger is 02-evidence-ledger-v04.csv.

Tables 1 and 2 are written by analysis/build_tables_v03.py and are pasted into this manuscript and its Chinese counterpart without hand editing; the generated files are results/table1-obligation-coverage-v03.en.md, results/table1-obligation-coverage-v03.zh.md, results/table2-adherence-patterns-v03.en.md and results/table2-adherence-patterns-v03.zh.md. The same script writes the dated timeline underlying Figure 1 to results/table3-timeline-v03.md, the parse of the advertised consequences to results/pathway-consequences-check-v04.md, and the intervals quoted in Sections 2, 6 and 8 to results/derived-intervals-v03.csv. The commitment and measure counts are recomputed from the three official chapter texts by analysis/extract_code_structure.py into results/code-structure.csv. Figure 1 is produced by analysis/make_figure1_calendar_v03.py and Figure 2 by analysis/make_figure2_mechanism_v03.py, each reading the input files named in its docstring from analysis/data/; each script is run once with –lang en and once with –lang

zh, and the English and Chinese renders come from the same script and the same data and differ only in the printed language. All paths in this section are relative to the directory containing this manuscript.

Statements

Funding. This work received institutional support from the Academy of Public Policy and Economics (APPE).

Competing interests. The authors are affiliated with APPE, which operates this journal. No external competing interests are declared.

References

Anthropic. (2025, July 21). Anthropic to sign the EU Code of Practice. https://www.anthropic.com/news/eu-code-practice.

CDT Europe. (2026, March 12). EU AI Act brief — Pt. 5, general-purpose AI models. Center for Democracy and Technology. https://cdt.org/insights/eu-ai-act-brief-pt-5-general-purpose-ai-models/.

Claburn, T. (2025, July 18). Meta declines to abide by voluntary EU AI safety guidelines. The Register. https://www.theregister.com/2025/07/18/meta_declines_eu_ai_guidelines/.

European AI Office. (2025). Code of Practice for General-Purpose AI Models: Signature form. European Commission. https://ec.europa.eu/newsroom/dae/redirection/document/118312.

European Artificial Intelligence Board. (2025, August 1). AI Board adequacy assessment: Code of Practice for general-purpose AI. European Commission. https://ec.europa.eu/newsroom/dae/redirection/document/118687.

European Commission. (2025a, July 17). AI Office invites providers to sign the GPAI Code of Practice. https://digital-strategy.ec.europa.eu/en/library/ai-office-invites-providers-sign-gpai-code-practice.

European Commission. (2025b, November 19). Commission guidelines on the scope of the obligations for providers of general-purpose AI models established by Regulation (EU) 2024/1689 (AI Act) (C(2025) 7719 final). https://ec.europa.eu/newsroom/dae/redirection/document/118340.

European Commission. (2025c, August 1). Commission opinion on the assessment of the General-Purpose AI Code of Practice within the meaning of Article 56 of Regulation (EU) 2024/1689 (C(2025) 5361 final). https://ec.europa.eu/newsroom/dae/redirection/document/118688.

European Commission. (2025d, December 5). Explanatory notice and template for the public summary of training content for general-purpose AI models required by Article 53(1)(d) of Regulation (EU) 2024/1689 (AI Act) (C(2025) 8311 final). https://digital-strategy.ec.europa.eu/en/library/explanatory-notice-and-template-public-summary-training-content-general-purpose-ai-models.

European Commission. (2025e, July 10). General-Purpose AI Code of Practice now available [Press release]. https://digital-strategy.ec.europa.eu/en/news/general-purpose-ai-code-practice-now-available.

European Commission. (2025f, September 9). General-purpose AI models in the AI Act: Questions and answers. https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers.

European Commission. (2025g, July 17). Signing the General-Purpose AI code of practice: Questions and answers. https://digital-strategy.ec.europa.eu/en/faqs/signing-general-purpose-ai-code-practice.

European Commission. (2026, January). Vademecum: The Signatory Taskforce of the General-Purpose AI Code of Practice. https://ec.europa.eu/newsroom/dae/redirection/document/124170.

Hoffmann, M. (2025, July 30). AI safety under the EU AI Code of Practice — A new global standard? Center for Security and Emerging Technology. https://cset.georgetown.edu/article/eu-ai-code-safety/.

Jones Day. (2025, August 6). EU AI Act: European Commission publishes General-Purpose AI Code of Practice. https://www.jonesday.com/en/insights/2025/08/eu-ai-act-european-commission-publishes-generalpurpose-ai-code-of-practice.

Larouche, P. (2025). Legal framework for an effective implementation of the AI Act (Issue Paper No. 2/2025). Centre on Regulation in Europe. https://cerre.eu/wp-content/uploads/2025/02/Legal-Framework-for-an-Effective-Implementation-of-the-AI-Act_FINAL.pdf.

Oliver, N., & Bommasani, R. (2025). Code of Practice for General-Purpose AI Models: Transparency chapter. European Commission. https://ec.europa.eu/newsroom/dae/redirection/document/118120.

OpenAI. (2025, July 11). The EU Code of Practice and future of AI in Europe. https://openai.com/global-affairs/eu-code-of-practice/.

Peukert, A. (2026). A new kid on the EU's co-regulatory block: Codes of practice under the AI Act. European Journal of Risk Regulation. Advance online publication. https://doi.org/10.1017/err.2026.10092.

Peukert, A., & Castets-Renard, C. (2025). Code of Practice for General-Purpose AI Models: Copyright chapter. European Commission. https://ec.europa.eu/newsroom/dae/redirection/document/118115.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024. http://data.europa.eu/eli/reg/2024/1689/oj.

Samwald, M., Bengio, Y., Schaake, M., Ziosi, M., Privitera, D., Reuel, A., Zacherl, A., Rajkumar, N., & Anderljung, M. (2025). Code of Practice for General-Purpose AI Models: Safety and Security chapter. European Commission. https://ec.europa.eu/newsroom/dae/redirection/document/118119.

Skadden, Arps, Slate, Meagher & Flom. (2025, August 14). EU's general-purpose AI obligations are now in force, with new guidance. https://www.skadden.com/insights/publications/2025/08/eus-general-purpose-ai-obligations.

Subin, S. (2025, July 18). Meta says it won't sign Europe AI agreement, calling it an overreach that will stunt growth. CNBC. https://www.cnbc.com/2025/07/18/meta-europe-ai-code.html.

Walker, K. (2025, July 30). We will sign the EU AI Code of Practice. Google. https://blog.google/around-the-globe/google-europe/eu-ai-code-practice/.