DEEPFAKES AND THE COMPLIANCE VACUUM: LABELLING AI-GENERATED CONTENT UNDER ARTICLE 50 OF THE EU AI ACT

Dr. Deepika Rani[1], Dr. Jay Kumar Pandey[2] and Dr. Rajendra Kumar[3]

ABSTRACT

Article 50 of the EU Artificial Intelligence Act requires providers to mark synthetic content "in a machine-readable format" so that it is "detectable as artificially generated or manipulated", and requires deployers to disclose deepfakes. It applies from 2 August 2026, yet it leaves its operative terms to "the generally acknowledged state of the art, as may be reflected in relevant technical standards". This article asks when AI-media duties actually bind, and who fixes what they require. It codes forty-seven obligations drawn from fifteen instruments in five jurisdictions at two snapshot dates, 1 January 2026 and 1 January 2028, with every commencement date verified against a primary text. Thirty obligations (63.8 per cent; 95% CI 49.5–76.0) were operative on 1 January 2026, and sixteen were enacted but not yet in application. Independently of the calendar, between five and fifteen of the thirty operative duties carry a residual constraint that commencement does not cure, depending on whether an author-coded "undefined term" flag, which an independent automated coder did not reproduce, is admitted. Doctrinally, the article shows that Article 50 combines delegation with no presumption of conformity: harmonised standards and common specifications under Articles 40 and 41 do not reach it, the legislature has stated that the codes of practice on which it relies "do not grant a presumption of conformity", and the Commission has said that adherence to the Code "does not constitute conclusive evidence of compliance", while telling supervisors to focus on whether signatories adhere to it. On the reasoning of *James Elliott*, a private specification cannot therefore acquire the status of EU law that a harmonised standard has. A paired contrast with China, where a mandatory labelling standard accompanied the duty, shows that this architecture is a design choice. The article closes with remedies available to the Commission and to regulated parties.

[1] Assistant Professor, School of Legal Studies, Jigyasa University, Dehradun, Uttarakhand, India. ORCID: 0009-0008-5569-4050.
[2] Assistant Professor, Department of Electrical & Electronics Engineering, Shri Ramswaroop Memorial University, Barabanki, Uttar Pradesh, India. ORCID: 0000-0003-4086-5730.
[3] Assistant Professor, School of Legal Studies, Babu Banarasi Das University, Lucknow, Uttar Pradesh, India. ORCID: 0000-0002-1194-3157.

Keywords: AI Act; Article 50; deepfakes; labelling of AI-generated content; labeling; content provenance; C2PA; harmonised standards; presumption of conformity; regulatory delay

I. Introduction

On 10 June 2026 the European Commission published the final Code of Practice on Transparency of AI-Generated Content. The Code is the instrument on which providers are expected to rely in meeting Article 50 of the Artificial Intelligence Act, which began to apply fewer than eight weeks later, on 2 August 2026.[4] On the question that matters most for synthetic media, namely what a compliant mark looks like, the Code is candid: "At the time of publication of this Code, relevant interoperability standards and/or best practices are yet to be developed, except for digitally signed metadata."[5] It names no specification for that metadata, and no provenance standard anywhere in its text.[6]

The Act itself explains why. Article 50(2) requires providers of generative systems to ensure that outputs "are marked in a machine-readable format and detectable as artificially generated or manipulated", and that their technical solutions are "effective, interoperable, robust and reliable as far as this is technically feasible, taking into account the specificities and limitations of various types of content, the costs of implementation and the generally acknowledged state of the art, as may be reflected in relevant technical standards".[7] The duty is binding; its content is delegated to the state of the art. The ordinary EU route from technical content to legal safe harbour, a harmonised standard whose reference is published in the Official Journal and which raises a presumption of conformity, does not extend to Article 50.[8] A provider that follows a standard, a specification or the Code therefore obtains no legal presumption that it has

[4] European Commission, Code of Practice on Transparency of AI-Generated Content (final text, June 10, 2026) [hereinafter Code of Practice]; Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence, OJ L, 2024/1689, 12.7.2024 [hereinafter AI Act], art. 113 (general date of application, 2 August 2026). The date of application of Chapter IV, which contains Article 50, is not altered by Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence, OJ L, 2026/1744, 24.7.2026 [hereinafter Digital Omnibus on AI], art. 1(40).

[5] Code of Practice, *supra* note 4, Measure 3.4 (Interoperability).

[6] *Id.* The text contains no reference to C2PA or to any named provider, and the only technical standards it names are ETSI EN 301 549 and WCAG 2.1, in an accessibility context. Sub-measure 1.1.1 requires "digitally signed metadata" without naming a format.

[7] AI Act, *supra* note 4, art. 50(2).

[8] AI Act, *supra* note 4, arts. 40(1), 41(1); Regulation (EU) No 1025/2012 of the European Parliament and of the Council of 25 October 2012 on European Standardisation, OJ L 316, 14.11.2012, p. 12, art. 2(1)(c) [hereinafter Regulation 1025/2012]. See *infra* Part VII.A.

complied, although, as Part VII.D shows, the Commission has told supervisors to focus on whether signatories adhere to the Code.

This article examines that arrangement from two directions. The first is temporal. A duty may be enacted without yet applying, and may apply without yet being performable, and an analysis that records one state per instrument cannot see the difference. The second is doctrinal. If the content of a binding duty is left to private technical work, the legal question is what status that work acquires, and the case law of the Court of Justice supplies a test. The article asks four questions.

RQ1. What share of a purposive sample of AI-media obligations was operative on 1 January 2026, and on 1 January 2028?

RQ2. Independently of the calendar, which residual constraints (an undefined operative term, a capability gap, or a delegation to another instrument) survive an obligation's commencement, and how reproducible is that coding?

RQ3. For Article 50, which actor fixes the operative content of the duty, and through what legal mechanism, if any, does that content acquire legal effect?

RQ4. Is the position of the EU's content-labelling duty a general feature of AI regulation, or a design choice that another jurisdiction made differently?

The article makes three contributions. It supplies a dataset in which every commencement date is verified against a primary text and recorded with a verbatim quotation. It proposes a two-axis codebook, separating temporal state from residual constraint, and reports how far an independent coder reproduces the second axis. And it analyses Article 50 as delegation without presumption, locating the provision within the doctrine of *James Elliott* and *Public.Resource.Org*.

Three claims that circulate in this field are not made. The article does not claim that the EU requires any particular provenance specification, such as C2PA; it does not. It does not claim that the Digital Omnibus deferred Article 50 because standards were late. The Omnibus left Article 50's date of application unchanged and added a transition, to 2 December 2026, for systems already on the market, for a stated reason of adaptation time; its recitals 2 and 40,

which attribute delay to standards, concern the high-risk obligations.[9] And it does not claim that the sample is representative: obligations were selected, not exhaustively enumerated, so the proportions describe this sample.

Part II situates the argument in the literature on private standard-setting and on AI Act standardisation. Part III sets out the concepts, and Part IV the method, including a validity test of the coding. Parts V and VI report the temporal and residual findings. Part VII analyses Article 50 as delegation without presumption, and contrasts it with China. Part VIII discusses implications and remedies, and Part IX concludes.

II. Literature and Theoretical Framework

A. Private standards as public law

EU standardisation rests on a division of labour that the standardisation bodies state plainly: "Standards are voluntary tools of the market whereas legislation is the prerogative of the public authorities and is mandatory in application."[10] The Court of Justice has complicated that division. In *James Elliott*, the Court held that a harmonised standard whose references have been published in the Official Journal "forms part of EU law, since it is by reference to the provisions of such a standard that it is established whether or not the presumption laid down [in the directive] applies to a given product".[11] The reasoning is conditional. Although the standard is developed by "an organisation governed by private law", it is "a necessary implementation measure which is strictly governed by the essential requirements defined by that directive, initiated, managed and monitored by the Commission, and its legal effects are subject to prior publication by the Commission of its references" in the Official Journal.[12] In *Public.Resource.Org*, the Grand Chamber, holding that the requested standards "form part of EU law owing to their legal effects", annulled the Commission's refusal of access, relying on the rule of law and the principle that individuals must be able to ascertain their rights and

[9] Digital Omnibus on AI, *supra* note 4, art. 1(39)(b) (inserting Article 111(4) of the AI Act), recital 38 (transitional period "to adapt their practices within a reasonable time without disrupting the market"); *id.* recitals 2, 40 (delay of standards and the high-risk obligations of Chapter III, Sections 1 to 3).
[10] CEN-CENELEC, Guide 30: European Guide on Standards and Regulation — Better Regulation through the Use of Voluntary Standards, Executive Summary (Edition 1, June 2015).
[11] Case C-613/14, James Elliott Construction Ltd v. Irish Asphalt Ltd, ¶ 40 (Oct. 27, 2016) (CELEX 62014CJ0613).
[12] *Id.* ¶ 43.

obligations.[13]

Commentators read these decisions as drawing technical standardisation into the Union legal order. Cuccuru argues that the Court extended its jurisdiction over harmonised standards, "paving the way for a deeper intersection between European judiciary and technical standardization";[14] Senden concludes that, as harmonised standards acquire "an increasingly EU public law nature", it "can be doubted whether the constitutional safeguards the Regulation provides are sufficient".[15] A further line of argument on transparency and the rule of law has been advanced by Davies and Van Waeyenberge.[16]

The point that matters here is the conditionality. On this case law, a private technical text has legal status because a mandate, a governing legal framework, Commission supervision and Official Journal publication are present. Whether a specification written outside that chain can claim the same status is a question this article returns to in Part VII.

B. Standardisation under the AI Act

Empirical work on AI Act standardisation concerns the high-risk requirements. Kilian, Jäck and Ebel interviewed twenty-three organisations developing AI applications. They report an effective implementation period of likely less than six months against at least twelve months actually required for around thirty technical standards, committees "predominantly influenced by large enterprises", and compliance costs of €100,000 to €300,000 a year. They also observe that if publication of the standards "is further delayed", then "the EU legislator could, as a last resort, postpone the application of the AI Act high-risk requirements".[17] Perarnaud notes that the responsible joint technical committee, CEN-CENELEC JTC 21, does not publicly share information on its membership or working process.[18] Hennessy and Mügge analyse the

[13] Case C-588/21 P, Public.Resource.Org, Inc. v. European Commission, ¶¶ 80–81, 89 (Mar. 5, 2024) (CELEX 62021CJ0588), setting aside Case T-185/19, Public.Resource.Org, Inc. v. European Commission (Gen. Ct. July 14, 2021).
[14] Pierluigi Cuccuru, *The Public and Private Sides of Harmonized Standards: James Elliott Construction v. Irish Asphalt*, 19 German L.J. 1399 (2018) (abstract and conclusion read; the body was not read in full).
[15] Linda Senden, *The Constitutional Fit of European Standardization Put to the Test*, 44 Legal Issues Econ. Integration 337 (2017) (abstract and introduction read).
[16] Zachariah Davies & Arnaud Van Waeyenberge, *Better Regulation by Standards? Harmonized Technical Standards, Transparency, and the Rule of Law*, 62 Common Mkt. L. Rev. 147 (2025). Cited for the existence of this line of argument only; the full text could not be obtained.
[17] Robert Kilian, Linda Jäck & Dominik Ebel, *European AI Standards – Technical Standardisation and Implementation Challenges under the EU AI Act*, 16 Eur. J. Risk Regul. 1038 (2025) (full text read; the figures and quotations are from the sections on key challenges and on structural barriers).
[18] Clément Perarnaud, *With the AI Act, We Need to Mind the Standards Gap* (CEPS, Apr. 25, 2023).

"delegation dilemma" the EU faces, strong incentives to delegate technical rulemaking and a need to retain control, and show that it deploys a continuum of hard and soft tactics; their analysis does not mention Article 50.[19]

C. Article 50 and the vagueness of its terms

Technical-legal work has begun to examine Article 50(2) directly. Schmitt and colleagues argue that compliance "cannot be reduced to post-hoc labeling" and identify structural gaps, including absent cross-platform marking formats and a misalignment between the "reliability" criterion and probabilistic model behaviour.[20] Souverain finds it difficult to translate the four criteria (reliable, interoperable, effective, robust) "into concrete and measurable evaluations" for language-model watermarks.[21] For the high-risk robustness requirement, Tamponi and colleagues describe the same structure: a horizontal Act that "leaves the task of providing technical guidance to harmonised standards".[22]

D. The gap

These studies establish that the operative terms are underspecified. They do not measure how widespread the condition is across AI-media duties, they do not separate it from the calendar, and they do not examine the legal status of whatever fills the space. This article does the first two empirically and the third doctrinally. It does not locate prior work that codes enactment against application across AI-media instruments, a statement about the searches run and not a proof of absence.[23]

III. Conceptual Foundations

An obligation passes through states that instrument-level analysis tends to merge. It is *drafted*

[19] Orla Hennessy & Daniel Mügge, *European Standardization in Flux: Navigating Delegation and Control in AI Governance*, 5 J. Standardisation (2026) (full text; a search of the text finds no mention of Article 50 or of watermarking).
[20] Vera Schmitt, Niklas Kruse, Premtim Sahitaj & Julius Schöning, *Transparency as Architecture: Structural Compliance Gaps in EU AI Act Article 50 II*, *in* Proceedings of the Language Resources and Evaluation Conference (LREC 2026) 1 (ELRA 2026), doi:10.63317/5hwpjzkvb4wf (abstract only).
[21] Thomas Souverain, *Watermarking Large Language Models in Europe: Interpreting the AI Act in Light of Technology*, 28 Ethics & Info. Tech. (2026) (abstract only).
[22] Roberta Tamponi, Carina Prunkl, Thomas Bäck & Anna V. Kononova, *Lost in Vagueness: Towards Context-Sensitive Standards for Robustness Assessment under the EU AI Act* (arXiv 2025) (abstract only).
[23] Searches of the OpenAlex index on 3 October 2026 combined terms for AI Act application dates and timelines, enacted-but-not-in-force obligations, harmonised standards and Article 50, regulatory indeterminacy, and content provenance with Article 50. Titles and abstracts returned were screened; none coded obligations by temporal state.

when proposed, *enacted* when adopted and in force as an act, *in application* when its provisions bind, and *performable* when the addressee can in fact comply. The gap between enactment and application is familiar and is cured by the calendar. The gap between application and performability is less familiar: a duty to operate a notice process is in application but cannot be performed before the law has created the process it presupposes.

Two different questions are therefore in play, and a single list of reasons a duty does not work conflates them. *Why does this duty not bind today?* is answered entirely by the calendar. *What will still be missing when it does bind?* is answered by the content of the duty and the capacity of its addressee, and the answer survives commencement. This article measures the two on independent axes. **Axis 1** records temporal state at a snapshot date. **Axis 2** records residual constraint, whether or not the duty is operative. An obligation in state E (enacted, not in application) is, by definition, one for which the calendar is the open question, so no finding is claimed from that identity. Independence matters for the converse case: an operative obligation can carry a residual constraint, and a single-axis coding of failure among the inoperative would never see it.

IV. Method

A. Sample and unit of analysis

The sample is purposive. It comprises forty-seven obligations from fifteen instruments in five jurisdictions (the European Union, the United Kingdom, China, India and the United States), chosen because they regulate synthetic media, platform handling of AI-generated content, likeness and personal data in AI media, or the AI systems that produce it, and because an official text could be obtained. The unit is the obligation, not the instrument, because obligations within one instrument reach operation on different dates. Obligations were selected, not exhaustively enumerated, and the number coded per instrument ranges from one to seven (Table 1). The proportions reported therefore describe this sample and should not be generalised to AI regulation.

Instrument	Jurisdiction	Obligations coded
Digital Services Act	EU	7
AI Act	EU	6
European Media Freedom Act	EU	5
Product Liability Directive	EU	2

Instrument	Jurisdiction	Obligations coded
China, Interim Measures (generative AI)	China	4
China, Measures for Labelling AI-Generated Content	China	4
Online Safety Act 2023	UK	4
Digital Personal Data Protection Act 2023	India	3
TAKE IT DOWN Act	US (federal)	3
18 U.S.C. § 1836	US (federal)	2
California AI Transparency Act (SB 942)	US (California)	3
California Elections Code § 20010	US (California)	1
California A.B. 2655	US (California)	1
Texas Election Code § 255.004	US (Texas)	1
Tennessee Public Chapter 588	US (Tennessee)	1
Total		**47**

Table 1. Instruments and obligations coded.

B. Axis 1: temporal state

Each obligation is assigned one of five states at each of two snapshot dates, 1 January 2026 and 1 January 2028.

Code	State	Definition
O	Operative	In application, and the addressee can presently comply and be lawfully penalised for breach
P	Partially operative	Some provisions in application, others deferred, delegated or unmet
E	Enacted, not in application	Adopted and in force as an act; obligations not yet applicable
N	Not adopted	Proposed or pending; no binding text
R	Repealed	Operative, then withdrawn by its own terms

Table 2. Temporal codebook (Axis 1).

The first snapshot is the first day of the year in which Article 50 becomes applicable; the second tests whether the position persists. Two snapshots show contrast, not a trend. Two rules govern state O. Under the *dischargeability rule*, an obligation is operative only where it is in

application and the addressee can in fact perform it. A duty is not demoted because compliance is costly, technically demanding or contested; where dischargeability is genuinely arguable the obligation is coded operative, so that the coding errs toward understating the vacuum. Under the *delegation rule*, a duty whose operative content is delegated to another instrument whose application is unverified is coded P. State R was added after a counterexample: the deceptive-media version of California Elections Code § 20010 is repealed by its own terms on 1 January 2027, when an older and different version becomes operative, which falsified an initial assumption that obligations only accrue.[24]

C. Axis 2: residual constraint

Independently of state, each obligation is coded for what remains unresolved. Three flags are recorded, and an obligation may carry more than one: **DEF**, an operative term is undefined by the instrument itself, so that compliance has no determinate content until someone construes it; **CAP**, the addressee cannot perform until an institutional or technical precondition exists; and **DEL**, the operative content is delegated to another instrument whose application is unverified. A flag is recorded only where the instrument itself leaves the matter open, not where compliance is merely difficult. The first author coded all forty-seven obligations and recorded the reasons row by row in the dataset.

D. Verification of dates

No date enters the dataset unless it is verified against a primary text and recorded with a verbatim quotation, a source URL and an access date. The commencement ledger holds 108 rows, 106 verified. The dataset-building script exits with an error if any obligation depends on an unverified date. Two items could not be verified and were excluded rather than coded: Michigan Compiled Laws § 168.932f, which could not be retrieved, and section 18 of the Online Safety Act 2023, which has never been brought into force and so has no commencement date. Verification changed the coding in several cases. A commencement instrument made in March 2025 was revoked before it operated;[25] the Indian data-protection Rules are dated 13,

[24] Cal. Elec. Code § 20010 (2026), as added by A.B. 730, 2019 Cal. Stat. ch. 493, and amended by A.B. 972, 2022 Cal. Stat. ch. 745, §§ 3–4 (the version added in 2019, which covers deceptive audio, video and images of candidates, remains in effect "only until January 1, 2027"; the version amended by § 4, a narrower rule confined to campaign material with superimposed pictures or photographs, becomes operative that day, so the coded duty ceases to apply and a different, narrower one replaces it). A.B. 502, 2026 Cal. Stat. ch. 333, amends only § 20012. Excluding the repealed duty raises the operative share at the second snapshot from 93.6 to 95.7 per cent.

[25] The Online Safety Act 2023 (Commencement No. 5) Regulations 2025, SI 2025/371 (UK), revoked by SI

not 14, November 2025;[26] and the Online Safety Act's reporting duty on child sexual exploitation and abuse content commenced on 7 April 2026, not in November 2025.[27] State statutory citations carry the year of the text consulted online (2026).

E. Validity: an independent second coder

Axis 2 rests on a judgement, and a single coder's judgement can be idiosyncratic. To measure that, an independent automated coder, TypeSafe's System One model Jev (version 1.13.0), was given each obligation blind: only the instrument name, the provision and the one-sentence obligation text, never the first coder's flags or reasons or the temporal states. For each flag it returned a probability, and a cut of 0.5 was used.[28] Table 3 reports agreement.

Flag	First coder	Second coder (≥0.5)	Agreement	Cohen's κ	AUC
DEF (as defined)	20/47	46/47	40.4%	−0.04	0.59
DEF (stricter wording)	20/47	38/47	57.4%	0.22	0.52
CAP	15/47	9/47	70.2%	0.23	0.77
DEL	1/47	12/47	76.6%	0.12	1.00

Table 3. Agreement between the first coder and the second coder on Axis 2. AUC is the probability that a randomly chosen first-coder positive receives a higher second-coder score than a randomly chosen negative.

The undefined-term flag was not reproduced. The second coder answered yes for nearly every duty, which is unsurprising, since all legal language has some open texture; a stricter wording that asked only for a material indeterminacy was written after seeing this result, is reported as a post-hoc sensitivity test, and did no better (AUC 0.52, chance level). The capability flag was reproduced moderately ($\kappa = 0.23$; AUC 0.77). The delegation flag has a single positive row, which the second coder ranked first, so its AUC says little. The second coder is a hosted model,

2025/1064, reg. 2, which records that the listed provisions "will no longer come into force on 3rd November 2025".

[26] The Digital Personal Data Protection Rules, 2025, G.S.R. 846(E) (India) (dated Nov. 13, 2025).

[27] The Online Safety Act 2023 (Commencement No. 7) Regulations 2026, SI 2026/262 (UK).

[28] The model and its question wording, the raw probabilities and the agreement script are supplied with the datasets. The first stricter wording was written after the first result was seen. TypeSafe, *System One API* (https://docs.typesafe.ai/api.md, consulted Oct. 3, 2026).

not a lawyer; it read a one-sentence paraphrase written by the first coder; and agreement therefore measures the replicability of the judgement given that paraphrase, not its correctness. The consequence is drawn in Part VI: findings that depend on DEF are reported as the upper end of a range, and the reproducible flags as the lower.

F. Limits

The sample is purposive and the obligations selected, so the proportions are descriptive of this sample. Two snapshots show contrast, not a trend. Dischargeability involves judgement. The second coder tested one aspect of Axis 2 only; Axis 1 is computed from verified dates and was not second-coded.

V. Finding I: The Vacuum Measured

On 1 January 2026, 30 of the 47 obligations were operative (63.8 per cent; 95% CI 49.5–76.0), one was partially operative, and sixteen were enacted but not yet in application (34.0 per cent; 22.2–48.3) (Table 4). Proportions are reported as exact fractions with Wilson score intervals; at n = 47 the intervals are wide, and that width is the honest description of what a sample of this size supports.

State	1 Jan 2026	1 Jan 2028
O, operative	30/47 (63.8%; 49.5–76.0)	44/47 (93.6%; 82.8–97.8)
P, partially operative	1/47 (2.1%)	2/47 (4.3%)
E, enacted, not in application	16/47 (34.0%; 22.2–48.3)	0/47
N, not adopted	0/47	0/47
R, repealed	0/47	1/47 (2.1%)

Table 4. Temporal state at both snapshots (95% Wilson intervals).

The sixteen obligations not yet in application at the first snapshot sat in the AI Act (five), California SB 942 (three), the Indian data-protection Act (three), the Product Liability Directive (two), the TAKE IT DOWN Act (two) and the Media Freedom Act (one). By 1 January 2028, forty-four obligations were operative and none was merely enacted. Two were partially operative: China's Interim Measures, article 12, whose labelling content is delegated to a separate instrument, and the AI Act's high-risk requirements, which apply to Annex III

systems from 2 December 2027 but to Annex I systems only from 2 August 2028.[29] One was repealed. Excluding the repealed duty, 44 of 46 (95.7 per cent; 85.5–98.8) were operative. The temporal vacuum is therefore real and closes within two years. That is a smaller claim than it first appears, as Part VI shows.

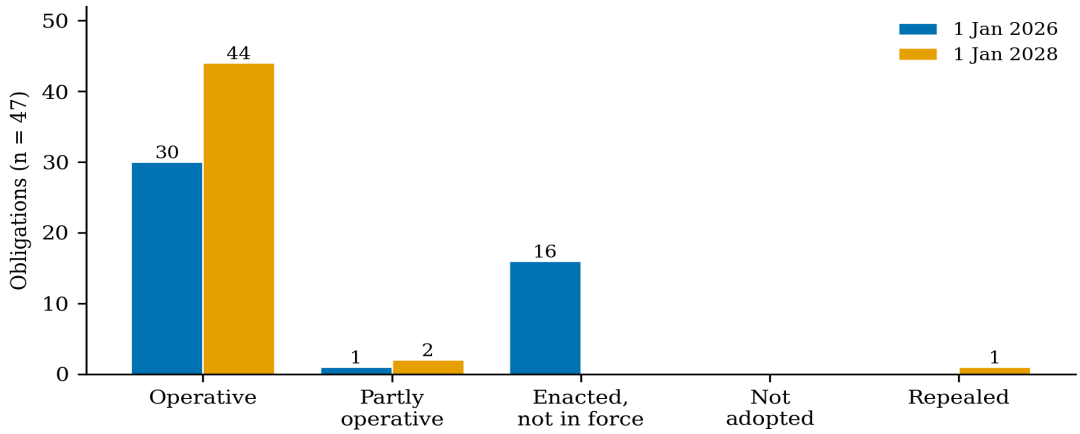


Figure 1. *Temporal state of the 47 coded obligations at the two snapshot dates. Computed from Live_Duty_Matrix.csv.*

Operativeness differed by regulatory object (Table 5), with conduct-side duties more often operative (24 of 33, 72.7 per cent) than content-side duties (6 of 12, 50.0 per cent). The difference is not statistically established at this sample size (Fisher's exact test, two-sided p = 0.17), the intervals overlap, and the two product-side obligations are too few to support any inference. The dischargeability rule moved the headline by a single obligation: a test that asked only whether a provision was in application would have counted 31 of 47 as operative. The rule matters for individual codings, such as the senior-manager offence in the Online Safety Act, and not for the aggregate.

Regulatory object	Operative, 1 Jan 2026	95% CI
Conduct	24/33 (72.7%)	55.8–84.9
Content	6/12 (50.0%)	25.4–74.6
Product	0/2 (0.0%)	0.0–65.8

Table 5. Operativeness by regulatory object.

[29] Digital Omnibus on AI, *supra* note 4, art. 1(40)(b) (replacing point (c) of the third paragraph of article 113 of the AI Act).

VI. Finding II: What Commencement Does Not Cure

The temporal axis fixes the size of the vacuum. The residual axis describes what is left once the calendar has run. Of the 47 obligations, 30 carry at least one flag (63.8 per cent; 95% CI 49.5–76.0): 20 an undefined operative term (42.6 per cent), 15 a capability gap (31.9 per cent) and one a delegation gap (Table 6). The flags co-occur, so the rows do not sum.

Residual constraint	Obligations	Share of 47 (95% CI)
DEF, undefined operative term	20/47	42.6% (29.5–56.7)
CAP, capability or build-out gap	15/47	31.9% (20.4–46.2)
DEL, delegated to an unverified instrument	1/47	2.1% (0.4–11.1)
Any flag	30/47	63.8% (49.5–76.0)

Table 6. Residual constraint across all 47 obligations (first-coder flags).

A. A duty can bind without saying what compliance requires

Restricting the residual axis to the thirty operative obligations gives a range, because the two sets of flags have different standing. Fifteen of the thirty (50.0 per cent; 33.2–66.8) carry at least one flag, thirteen of them an undefined operative term. Five (16.7 per cent; 7.3–33.6) carry a flag that the second coder at least partly reproduced, a capability or delegation gap. Fifteen carry no flag (Table 7). The defensible statement is therefore that between a sixth and a half of the operative obligations in this sample bind their addressee while leaving something unresolved that commencement will not cure, and that the upper figure rests on a flag whose reproducibility the second coder did not support.

Among the 30 operative obligations	Obligations	Share (95% CI)
Any flag (upper bound; includes DEF)	15/30	50.0% (33.2–66.8)
Capability or delegation flag only (reproducible subset)	5/30	16.7% (7.3–33.6)
No flag	15/30	50.0% (33.2–66.8)

Table 7. Residual constraint among operative obligations, with the reproducible subset.

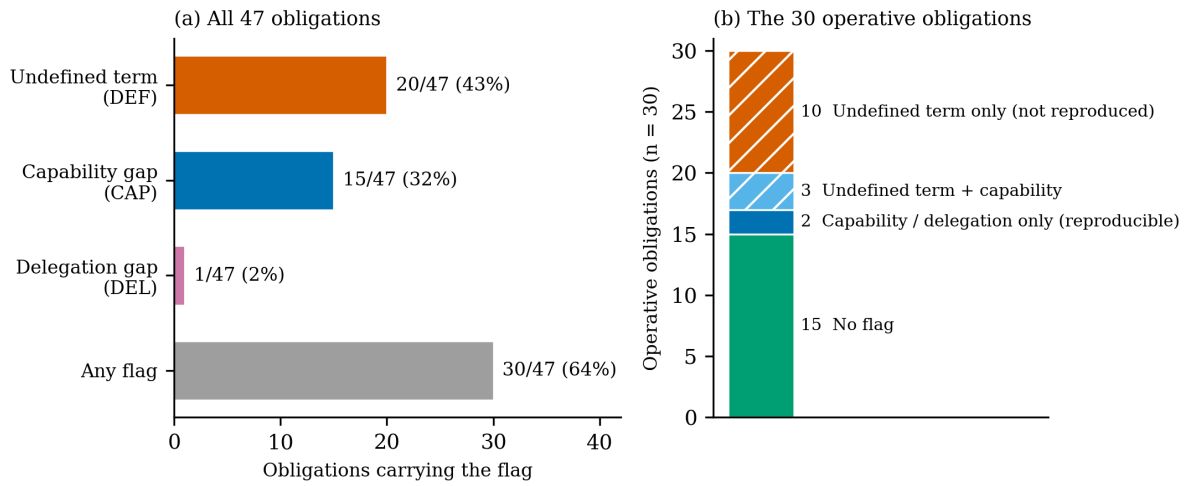


Figure 2. *Residual constraint. (a) Obligations carrying each flag, across all 47. (b) The 30 operative obligations by the flags they carry; the undefined-term flag is shown separately because it was not independently reproduced. Computed from Live_Duty_Matrix.csv.*

The single-axis alternative illustrates why the separation matters. A study that coded failure only among the inoperative would look for undefined terms where the inoperative obligations are, and would miss those attached to obligations that already apply. Article 50(2) is the clearest example in the sample: it is enacted but not yet applicable on the first snapshot date, and when it applies it will carry an undefined term ("machine-readable", "detectable") and a capability gap (no settled marking or detection specification), only one of which the calendar cures.[30] In the other direction, Article 37 of the Digital Services Act is operative and carries a capability gap, since it presupposes an accredited audit market; the Online Safety Act's senior-manager offence is operative with no flag.

B. Patience resolves few of the not-yet-applicable duties

Seventeen obligations were not operative on 1 January 2026, sixteen enacted and one partial. For a not-yet-applicable duty the calendar is the visible question, but not the only one. Of the sixteen, fourteen carry at least one flag (87.5 per cent), and ten carry the reproducible capability or delegation flag (62.5 per cent; 38.6–81.5). Across the seventeen, the date alone resolves between two (no flag of any kind) and six (no reproducible flag) obligations. An organisation that reads *not yet applicable* and waits has read its position correctly for a minority of them.

[30] AI Act, *supra* note 4, art. 50(2); Code of Practice, *supra* note 4, Measure 3.4.

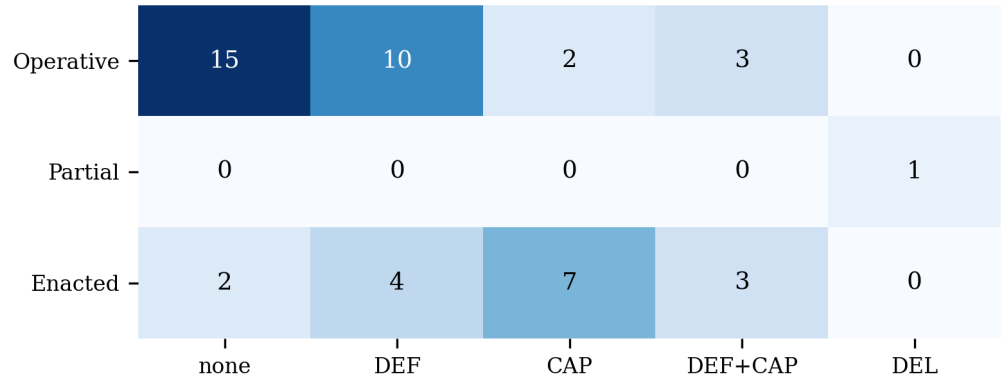


Figure 3. *Temporal state against residual constraint at 1 January 2026. Computed from Live_Duty_Matrix.csv.*

VII. Finding III: Article 50 as Delegation without Presumption

A. What Article 50(2) delegates

Article 50(2) states the duty and then points outward for its content. Providers must ensure their solutions are "effective, interoperable, robust and reliable as far as this is technically feasible", measured against "the generally acknowledged state of the art, as may be reflected in relevant technical standards". The Act does not define "machine-readable", "detectable", "synthetic" or "watermark", and Article 3 supplies no definition of the "state of the art" that governs; recital 133 lists, as examples, "watermarks, metadata identifications, cryptographic methods for proving provenance and authenticity of content, logging methods, fingerprints or other techniques, as may be appropriate".[31] The operative content is thus left to a moving external benchmark.

B. Why no standard can confer a presumption

The EU's usual mechanism for converting technical content into legal safe harbour is a harmonised standard, which Regulation 1025/2012 defines as a European standard "adopted on the basis of a request made by the Commission for the application of Union harmonisation legislation".[32] Under the AI Act, the presumption of conformity in Article 40(1) attaches to

[31] AI Act, *supra* note 4, arts. 3, 50(2), recital 133. Article 3 defines "deep fake" (art. 3(60)) but contains no definition of "machine-readable", "detectable", "synthetic", "watermark", "editorial" or "state of the art".
[32] Regulation 1025/2012, *supra* note 8, art. 2(1)(c); *see also id.* art. 10(1) (standardisation requests; standards to be "market-driven").

compliance with harmonised standards whose references have been published in the Official Journal, but only for "the requirements set out in Section 2" of Chapter III (high-risk systems) and the obligations in Sections 2 and 3 of Chapter V (general-purpose models). The Commission's duty to issue standardisation requests in Article 40(2), and its power to adopt common specifications in Article 41(1), are confined in the same way. Article 50 sits in Chapter IV. No harmonised standard, and no common specification, can raise a presumption that a provider has complied with it (Table 8).

	High-risk requirements (Chapter III, Section 2)	Article 50(2) marking duty
Operative content fixed by	Harmonised standards; common specifications	"State of the art, as may be reflected in relevant technical standards"
Presumption of conformity	Yes (art. 40(1))	None
Standardisation request	Required (art. 40(2))	None provided for
Common specifications fallback	Yes (art. 41)	None
Effect of the Digital Omnibus on the date	Deferred to 2 Dec. 2027 (Annex III) and 2 Aug. 2028 (Annex I), citing late standards	Unchanged (2 Aug. 2026); four-month transition for systems already on the market
Soft-law vehicle	Not applicable	Code of Practice (art. 50(7)): "limited legal effect"

Table 8. Two routes by which the AI Act turns technical content into legal effect.[33]

This explains a negative result that might otherwise look like failure. The deliverables listed on the working-groups page of CEN-CENELEC JTC 21 concern risk management, quality management, trustworthiness, dataset quality, bias, logging, conformity assessment and cybersecurity, among others; none concerns the marking or labelling of synthetic content.[34] That does not mean no standard is expected. The Commission's final guidelines anticipate that interoperable detection will depend, for a time, on providers' own or shared solutions "until harmonised standards and a standardised provider-agnostic interoperable detection solution emerge".[35] Under Article 40(1), however, such a standard would raise no presumption of

[33] AI Act, *supra* note 4, arts. 40(1)–(2), 41(1), 50(7), 113; Digital Omnibus on AI, *supra* note 4, art. 1(20), (39)(b), (40)(b), recitals 2, 38, 40, 41.
[34] CEN-CENELEC JTC 21, *Working Groups and Deliverables* (https://jtc21.eu/working-groups/, consulted Oct. 3, 2026).
[35] Final Guidelines, *supra* note 40, ¶ 76.

conformity with Article 50. The point also puts the standards channel in context: even for the high-risk requirements, where the mechanism exists, no AI Act harmonised standard has been cited in the Official Journal, and the furthest any has advanced is prEN 18286 (quality management), which entered public enquiry on 30 October 2025.[36]

C. The *James Elliott* test applied

On the case law set out in Part II, a private technical text is part of EU law where several conditions are present: it is developed on a Commission mandate; it is "strictly governed by the essential requirements" of the legislation it implements; it is "initiated, managed and monitored by the Commission"; and its references have been published in the Official Journal.[37] None of these is available for a specification written to meet Article 50(2). It follows, as an inference from the Court's reasoning and not as a holding, since the Court did not address Article 50, that a CEN, ISO or consortium specification for marking synthetic content would not "form part of EU law". It would be evidence of the state of the art, which Article 50(2) makes the benchmark, but not a safe harbour. A provider's reliance on any such text carries no legal presumption. The rule-of-law concern in *Public.Resource.Org*, that individuals "must be able to ascertain unequivocally what their rights and obligations are", is engaged here in a different form. Access to the text is not the issue, because the text is not law. The issue is that the duty binds, its content is set by texts that are neither mandated nor published by the Commission, and an addressee cannot ascertain from the law what compliance requires.

D. The soft-law vehicle and its legal effect

The Act provides one route to guidance: codes of practice under Article 50(7). The Digital Omnibus reshaped it. It removed the Commission's power to approve such codes by implementing act, replaced it with a duty to "assess whether adherence to those codes of practice is adequate to ensure compliance with the obligations laid down in paragraphs 2 and 4", and retained a power to adopt an implementing act "specifying common rules for the implementation of those obligations" if a code is inadequate.[38] Its recital explains the choice:

[36] European Commission, *AI Act Standardisation* (https://digital-strategy.ec.europa.eu/en/policies/ai-act-standardisation, consulted Oct. 3, 2026) (describing prEN 18286 as "the first harmonised standard for AI to enter public enquiry"); European Commission, list of implementing decisions on harmonised standards (swept in full; none concerns the AI Act).
[37] James Elliott, *supra* note 11, ¶¶ 40, 43; Public.Resource.Org, *supra* note 13, ¶ 81 (quoting the rule-of-law principle).
[38] Digital Omnibus on AI, *supra* note 4, art. 1(20) (replacing art. 50(7) of the AI Act), recital 41.

codes of practice "have limited legal effect, and in particular do not grant a presumption of conformity". The Commission must also issue guidelines on the practical implementation of Article 50, which must "take due account of the generally acknowledged state of the art" and of harmonised standards and common specifications referred to in Articles 40 and 41,[39] instruments that, as Part VII.B showed, do not exist for Article 50.

The Commission's final guidelines do supply interpretation. They understand the "state of the art" as "a developed stage of technical capability at a given time as regards products, processes and services, based on the relevant consolidated findings of science, technology and experience and which is accepted as good practice in technology", a formulation they source to the annex of the Commission's standardisation request for high-risk systems, and they define "human review" as "the deliberate examination of the substance of the content by one or more natural persons possessing relevant knowledge and professional judgement", with fact-checking "a minimum requirement".[40] That interpretation narrows the practical indeterminacy of some terms. It does not change their legal character: the guidelines state that they are "non-binding" and that "[a]ny authoritative interpretation of the AI Act may ultimately only be given by the Court of Justice".[41]

What the Commission has done with the Code matters more than that interpretation, and it refines the claim of this Part. On 8 July 2026 the Commission concluded that the Code "adequately covers the obligations provided for in Articles 50(2), (4) and (5) AI Act" and that providers "may rely on the Code", "while recognising that adherence to the Code does not constitute conclusive evidence of compliance with these obligations".[42] The guidelines then describe adherence as "a straightforward, predictable, and legally certain way of demonstrating compliance", state that supervisors "will focus their supervisory activities on assessing whether those signatories have adhered to the code", treat commitments under the Code as a possible mitigating factor in fixing fines, and warn that authorities "will likely need more detailed

[39] AI Act, *supra* note 4, art. 96(1)(d) and the last subparagraph of art. 96(1).
[40] European Commission, Communication: Guidelines on the Implementation of the Transparency Obligations for Certain AI Systems under Article 50 of Regulation (EU) 2024/1689, C(2026) 5054 final (July 20, 2026), annex, ¶¶ 83, 134 & n.28 [hereinafter Final Guidelines]. The footnoted source for the definition of the state of the art is Annex II of Commission Implementing Decision C(2025) 3871 final on a standardisation request in support of the AI Act.
[41] Final Guidelines, *supra* note 40, ¶ 5.
[42] Commission Opinion on the Assessment of the Code of Practice on Transparency of AI-Generated Content in Accordance with Articles 50(7) and 56(6) of Regulation (EU) 2024/1689, C(2026) 4839 final (July 8, 2026), ¶ 52. The AI Board adopted its adequacy assessment on July 9, 2026.

information" from non-signatories.[43] The Union therefore has delegation without a *legal* presumption, partly compensated by a *supervisory* benefit. The distinction is significant for a regulated party. A supervisory benefit shapes how an authority treats a signatory. It does not bind a court, and it does not settle the construction of an open term in litigation between private parties or on an appeal against a fine.

The Code of Practice that resulted defers to the missing technical content. It records that "relevant interoperability standards and/or best practices are yet to be developed, except for digitally signed metadata", and its signatories "commit to implement those standards in a timely manner as they emerge".[44] The legislature has thus stated, in a recital, the limited legal effect of the instrument on which the Commission relies, and that instrument in turn defers to specifications that do not yet exist.

E. What filled the space

Technical content was nonetheless being produced, by a private consortium and ahead of the final Code. In April 2026, three months before Article 50 began to apply, version 2.4 of the C2PA specification introduced an "AI Disclosure" assertion "to provide a means for a Claim Generator to provide machine-readable AI transparency information", enabling "verifiable AI transparency, automated compliance verification, and trustworthy AI content at scale", and providing that it "shall have a label of c2pa.ai-disclosure".[45] The specification defines an "Invisible Watermark" and a "Manifest Repository", which together cover the two layers, signed metadata and imperceptible watermarking, that the Code's measures call for. It also supplies a three-value enumeration of human oversight (fully_autonomous, prompt_guided and human_validated).

That enumeration is relevant to an undefined statutory term. Article 50(4) already makes the deployer's duty to disclose AI-generated text turn on whether it "has undergone a process of human review or editorial control" and whether a person "holds editorial responsibility"; the Act defines neither phrase, although the Commission's final guidelines define human review

[43] Final Guidelines, *supra* note 40, ¶¶ 146–150 (quoting ¶¶ 147–148).
[44] Code of Practice, *supra* note 4, Measure 3.4.
[45] Coalition for Content Provenance and Authenticity, C2PA Technical Specification v. 2.4, §§ 2.4.3, 2.4.5, 5.3.1, 18.28.1–18.28.3 (Apr. 2026) (https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html, consulted Oct. 3, 2026). Version 2.4 is dated April 2026 in § 5.3.1; the downloadable bundle on the consortium's site still lists version 2.3.

and editorial control.[46] The specification's oversight enumeration is a second, independent candidate vocabulary for the same question. It is not a mapped one: a direct count of the full text of version 2.4 finds no occurrence of "AI Act" or of "2024/1689", and the Code of Practice does not name C2PA. The relationship is therefore best described as anticipatory and parallel, not as implementation or capture. A private specification and a legislative duty reached a comparable vocabulary by independent routes, and the law supplies no mechanism for saying which governs.

F. A paired contrast: China

The same technical question was answered differently in China. The Measures for Labeling AI-Generated Synthetic Content, issued by four state bodies in a notice dated 7 March 2025, took effect on 1 September 2025.[47] The mandatory national standard that supplies the technical method, GB 45438-2025, was published on 28 February 2025, before the Measures were dated, and took effect on the same day, 1 September 2025. The standard was legislated into force alongside the duty. The EU's Article 50(2) took the opposite route. The contrast is between two architectures, not a controlled comparison. The two systems differ in their rulemaking procedures and in how standards relate to law, and two cases cannot isolate a cause. What the pair shows is that the absence of a mandated technical specification at the date a labelling duty applies is a design choice, not a necessary feature of regulating synthetic media.

G. An asymmetry in the Omnibus

The Omnibus offers one further observation, stated without attributing intent. Its recital 2 records that "the delayed preparation of standards" for high-risk systems had produced "a compliance burden that is heavier than expected", and its recital 40 that "the delayed availability of standards, common specifications, and alternative guidance" did "not justify maintaining" the initial date of application "[f]or the obligations related to high-risk AI systems". The high-risk dates were moved accordingly, as Kilian and colleagues had anticipated as a last resort.[48] Article 50, which has no standards pathway at all, kept its original

[46] AI Act, *supra* note 4, art. 50(4), second subparagraph; *id.* art. 3; Final Guidelines, *supra* note 40, ¶ 134.
[47] Measures for Labeling AI-Generated Synthetic Content (人工智能生成合成内容□□□法), Guoxinban Tongzi [2025] No. 2 (China) (text dated Mar. 7, 2025; effective Sept. 1, 2025); GB 45438-2025, Cybersecurity Technology — Labelling Method for Content Generated by Artificial Intelligence (China) (published Feb. 28, 2025; mandatory; effective Sept. 1, 2025).
[48] Kilian, Jäck & Ebel, *supra* note 17 (the passage quoted in Part II.B).

date and received a four-month transition for systems already on the market, for a stated reason of adaptation time.[49] The obligation with a presumption mechanism was delayed because its standards were late; the obligation with none was not. These two recitals are therefore evidence that the legislature treats the availability of standards as the pacing item for the high-risk obligations. They are not, as an earlier version of this analysis assumed, evidence about Article 50.

VIII. Discussion

A. What the findings establish, and what they do not

The three findings rest on different foundations. The temporal result is the most secure. It is computed from verified commencement dates, and it is insensitive to the main coding rule: asking only whether a provision was in application would have changed the headline by a single obligation. The residual result is less secure. A range, from a sixth to a half of operative duties carrying a constraint that commencement does not cure, is the defensible statement, because the upper bound depends on a flag that an independent coder did not reproduce. The doctrinal result does not depend on the coding at all. It follows from the text of Articles 40, 41, 50 and 113 of the AI Act, from recital 41 of the Omnibus, and from the reasoning in *James Elliott*; its one inferential step, that a specification written outside the mandated chain would not form part of EU law, is identified as an inference.

Three limits on interpretation should be stated. The sample is purposive, so no proportion is a statement about AI regulation generally. The article does not attribute intent: the contrast between the deferred high-risk obligations and the undeferred Article 50 is an observation about the face of the Omnibus. And the China comparison is a paired contrast between two architectures, not a test of cause.

B. Implications for providers and deployers

Delegation without presumption has practical consequences. First, following a standard, a specification or the Code of Practice does not secure a presumption of compliance with Article 50(2). Such texts are evidence of the state of the art, which is the statutory benchmark, and a provider should treat them as evidence to be assembled and documented, not as a safe harbour.

[49] Digital Omnibus on AI, *supra* note 4, recitals 2, 38, 40, art. 1(39)(b), (40)(b).

Adherence to the Code earns something narrower, supervisory focus and a possible mitigating factor, which the Commission itself describes as not "conclusive evidence" of compliance. Second, where the Act leaves a term open, as it does for "machine-readable", "detectable" and, in Article 50(4), "editorial control", a provider or deployer should record the construction it adopted and the sources relied on, including the Commission's non-binding guidance, because only the Court of Justice can settle the meaning. Third, the capability gaps recorded for Article 50 are the part of the problem that preparation can address before the date of application, and the undefined terms are the part it cannot.

C. Remedies available to the institutions

The remedies differ in what legal consequence they can carry.

First, the adequacy assessment under Article 50(7) has now been made, and the Commission has said what adherence earns in supervisory practice. What it has not said, and cannot say by non-binding guidance, is what legal weight adherence carries before a court. Recital 41 says codes grant no presumption of conformity; the opinion says adherence is not conclusive evidence. A binding statement of the evidential weight of adherence would need an act of the legislature or an implementing act.

Second, if a code is judged inadequate, Article 50(7) permits an implementing act "specifying common rules for the implementation of those obligations". That is the only route in the Act by which the content of Article 50 can acquire binding specificity through the Commission.

Third, binding definitions of the open terms require an amendment of Article 3 or an implementing act. Guidelines, however useful, interpret and do not enact.

Fourth, as to standardisation, Regulation 1025/2012 permits the Commission to request European standards and deliverables in general terms.[50] Without extending the presumption in Article 40(1) to Chapter IV, however, a standard produced in response would not raise a presumption of conformity. Whether a Chapter IV standard could qualify as a "harmonised standard" under the definition in the Regulation, and whether the legislature should extend

[50] Regulation 1025/2012, *supra* note 8, art. 10(1) (the Commission "may within the limitations of the competences laid down in the Treaties, request one or several European standardisation organisations to draft a European standard or European standardisation deliverable within a set deadline").

Article 40(1), are questions for the legislature that this article does not resolve.

Fifth, the Chinese sequence suggests a design principle: where a technical specification is intended to carry legal effect, it should exist, and be published, on the date the duty applies.

D. Limitations

Four limitations bear on the conclusions. The sample is purposive and the obligations selected. The second coder tested one aspect of one axis, and the undefined-term flag was not reproduced. Several works in the literature could be read only at abstract level or not at all, as footnoted, and are cited for the existence of an argument only. And the second-coder test and the temporal coding were completed before the Commission's final guidelines and opinion were read, so the paraphrases coded for Article 50 reflect the statute and not those documents. The guidelines and opinion were then read in full and are reflected in Part VII.D. None of these limits affects the temporal findings, which are computed from verified dates, or the doctrinal analysis, which rests on the instruments' text.

IX. Conclusion

Thirty of forty-seven AI-media obligations in this sample were operative at the start of 2026, and by 2028 forty-four were. The temporal vacuum is real and closes within two years. A more durable problem sits beneath it: between a sixth and a half of the operative duties carry a constraint that the calendar does not cure, the upper figure resting on a judgement that an independent coder did not reproduce.

For the provision that matters most to synthetic media, the problem is structural. Article 50(2) binds providers to a result, marking that is machine-readable and detectable, and measures the technical solution against the state of the art "as may be reflected in relevant technical standards". The EU's ordinary mechanism for converting technical content into safe harbour does not reach it: Articles 40 and 41 are confined to the high-risk and general-purpose chapters, and the legislature has stated that the codes of practice on which Article 50 relies "do not grant a presumption of conformity". On the reasoning of *James Elliott*, a specification written outside the mandated and published chain does not form part of EU law. The result is delegation without presumption: a binding duty whose content is set by a moving benchmark, with interpretation from the Commission that is non-binding by its own account, a Code that

supervisors will treat as a benchmark but that is not conclusive evidence of compliance, and with a private specification, C2PA 2.4, that reached a comparable vocabulary three months before the duty applied without referring to it.

This is a design, and designs can be changed. China placed a mandatory standard in force on the day its labelling duty began. The Omnibus gave high-risk obligations more time, citing among other things the late availability of standards, and left Article 50 on its original date. The remedies are modest and lie within the Commission's existing powers: say what adherence to an adequate code earns, use the implementing-act power if a code falls short, and make binding the definitions that guidance can only suggest. Until then, a provider that complies with Article 50 does so on its own construction of terms the law has not settled, and carries the risk that the construction is wrong.

References

European Union instruments and decisions

1. 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.

2. Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 Amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as Regards the Simplification of the Implementation of Harmonised Rules on Artificial Intelligence (Digital Omnibus on AI). OJ L, 2026/1744, 24.7.2026.

3. Regulation (EU) No 1025/2012 of the European Parliament and of the Council of 25 October 2012 on European Standardisation. OJ L 316, 14.11.2012, p. 12.

4. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market for Digital Services and Amending Directive 2000/31/EC (Digital Services Act). OJ L 277, 27.10.2022, p. 1.

5. Regulation (EU) 2024/1083 of the European Parliament and of the Council of 11 April 2024 Establishing a Common Framework for Media Services in the Internal Market and Amending Directive 2010/13/EU (European Media Freedom Act). OJ L, 2024/1083, 17.4.2024.

6. Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on Liability for Defective Products and Repealing Council Directive 85/374/EEC. OJ L, 2024/2853, 18.11.2024.

7. Case C-613/14, *James Elliott Construction Ltd v. Irish Asphalt Ltd*, judgment of 27 October 2016 (CELEX 62014CJ0613).

8. Case C-588/21 P, *Public.Resource.Org, Inc. v. European Commission*, judgment of 5 March 2024 (CELEX 62021CJ0588), setting aside Case T-185/19, EU:T:2021:445.

9. European Commission. *Code of Practice on Transparency of AI-Generated Content*. Final text, 10 June 2026.

10. European Commission. *Guidelines on the Implementation of the Transparency Obligations for Certain AI Systems under Article 50 of Regulation (EU) 2024/1689*. C(2026) 5054 final, 20 July 2026; *Commission Opinion on the Assessment of the Code of Practice on Transparency of AI-Generated Content*. C(2026) 4839 final, 8 July 2026.

11. European Commission. *AI Act Standardisation*. digital-strategy.ec.europa.eu/en/policies/ai-act-standardisation. Accessed 3 Oct. 2026.

Instruments from other jurisdictions coded in the dataset

12. Online Safety Act 2023, c. 50 (UK); Online Safety Act 2023 (Commencement No. 1) Regulations 2023, SI 2023/1420; Online Safety Act 2023 (Commencement No. 5) Regulations 2025, SI 2025/371, revoked by SI 2025/1064; Online Safety Act 2023 (Commencement No. 7) Regulations 2026, SI 2026/262; Data (Use and Access) Act 2025 (UK) and SI 2025/982.

13. The Digital Personal Data Protection Act, 2023, No. 22 of 2023 (India); The Digital Personal Data Protection Rules, 2025, G.S.R. 846(E) (13 Nov. 2025); G.S.R. 843(E) (commencement notification).

14. Interim Measures for the Management of Generative Artificial Intelligence Services (生成式人工智能服务管理暂行办法), Order No. 15 of the Cyberspace Administration of China and six other bodies (10 July 2023; effective 15 Aug. 2023) (China).

15. Measures for Labeling AI-Generated Synthetic Content (人工智能生成合成内容标识办法), Guoxinban Tongzi [2025] No. 2 (dated 7 Mar. 2025; effective 1 Sept. 2025) (China).

16. GB 45438-2025, Cybersecurity Technology: Labelling Method for Content Generated by Artificial Intelligence (published 28 Feb. 2025; mandatory; effective 1 Sept. 2025) (China).

17. TAKE IT DOWN Act, Pub. L. No. 119-12 (2025) (United States); Defend Trade Secrets Act of 2016, Pub. L. No. 114-153, codified at 18 U.S.C. § 1836.

18. Ensuring Likeness, Voice, and Image Security Act of 2024, 2024 Tenn. Pub. Acts ch. 588, amending Tenn. Code Ann. §§ 47-25-1101 to -1107.

19. Cal. Elec. Code § 20010 (2026) (added by A.B. 730, 2019 Cal. Stat. ch. 493; amended by A.B. 972, 2022 Cal. Stat. ch. 745); Cal. Elec. Code § 20513 (2026) (added by A.B. 2655, 2024 Cal. Stat. ch. 261); California AI Transparency Act, S.B. 942, 2024 Cal. Stat. ch. 291, as amended by A.B. 853, 2025 Cal. Stat. ch. 674, Cal. Bus. & Prof. Code § 22757 *et seq.* (2026); Tex. Elec. Code Ann. § 255.004(d)–(e) (2026) (added by S.B. 751, 2019 Tex. Gen. Laws ch. 1339).

Technical specifications and institutional materials

20. Coalition for Content Provenance and Authenticity. *C2PA Technical Specification*, version 2.4 (Apr. 2026). spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html. Accessed 3 Oct. 2026.

21. CEN-CENELEC. *Guide 30: European Guide on Standards and Regulation — Better Regulation through the Use of Voluntary Standards*. Edition 1, June 2015.

22. CEN/CENELEC JTC 21. *Working Groups and Deliverables*. jtc21.eu/working-groups. Accessed 3 Oct. 2026.

23. TypeSafe. *System One API*. docs.typesafe.ai/api.md. Accessed 3 Oct. 2026.

Scholarship

24. Cuccuru, Pierluigi. "The Public and Private Sides of Harmonized Standards: James Elliott Construction v. Irish Asphalt." *German Law Journal*, vol. 19, no. 6, 2018, pp. 1399–1416. doi:10.1017/s2071832200023087. [abstract, introduction and conclusion read]

25. Davies, Zachariah, and Arnaud Van Waeyenberge. "Better Regulation by Standards? Harmonized Technical Standards, Transparency, and the Rule of Law." *Common Market Law Review*, vol. 62, no. 1, 2025, pp. 147–74. doi:10.54648/cola2025006. [not read]

26. Hennessy, Orla, and Daniel Mügge. "European Standardization in Flux: Navigating Delegation and Control in AI Governance." *Journal of Standardisation*, vol. 5, 2026. doi:10.59490/jos.2026.8346. [full text]

27. Kilian, Robert, Linda Jäck, and Dominik Ebel. "European AI Standards – Technical Standardisation and Implementation Challenges under the EU AI Act." *European Journal of Risk Regulation*, vol. 16, no. 3, 2025, pp. 1038–62. doi:10.1017/err.2025.10032. [full text]

28. Perarnaud, Clément. "With the AI Act, We Need to Mind the Standards Gap." Centre for European Policy Studies, 25 Apr. 2023.

29. Schmitt, Vera, Niklas Kruse, Premtim Sahitaj, and Julius Schöning. "Transparency as Architecture: Structural Compliance Gaps in EU AI Act Article 50 II." *Proceedings of the Language Resources and Evaluation Conference (LREC 2026)*, European Language Resources Association, 2026, pp. 1–11. doi:10.63317/5hwpjzkvb4wf. [abstract only]

30. Senden, Linda. "The Constitutional Fit of European Standardization Put to the Test." *Legal Issues of Economic Integration*, vol. 44, no. 4, 2017, pp. 337–52. doi:10.54648/leie2017018. [abstract and introduction read]

31. Souverain, Thomas. "Watermarking Large Language Models in Europe: Interpreting the AI Act in Light of Technology." *Ethics and Information Technology*, vol. 28, no. 4, 2026. doi:10.1007/s10676-026-09918-w. [abstract only]

32. Tamponi, Roberta, Carina Prunkl, Thomas Bäck, and Anna V. Kononova. "Lost in Vagueness: Towards Context-Sensitive Standards for Robustness Assessment under the EU AI Act." arXiv, 2025. doi:10.48550/arxiv.2511.15620. [abstract only]