



Comments on _____ Draft Guidelines on the Implementation of the Transparency obligations for certain AI systems under article 50 of regulation (EU) 2024/1689 (THE 'AI ACT')





A. Table of Contents

A. TABLE OF CONTENTS	2
B. LIST OF ABBREVIATIONS & DEFINITIONS	3
C. ABOUT THE CYRIL SHROFF CENTRE FOR AI, LAW & REGULATION	4
D. COMMENTS SUBMITTED TO THE EUROPEAN COMMISSION	5
E. DETAILED COMMENTS	8
1. The“Obviousness” Exception and the Average-Person Standard	8-9
2. Attenuated Disclosure for Artistic, Satirical or Fictional Deep Fakes	9-10
3. The Editorial Responsibility Exception for AI-Generated Text	10
4. Interplay with the General-Purpose AI Model Layer	10-11
5. Extra-Territorial Applicability and the Brussels Effect	11-12
F. CONCLUSION	13

B. List of Abbreviations & Definitions

Abbreviation / Term	Definition
AI	Artificial Intelligence as defined under Article 3(1) of the AI Act, i.e. “a machine-based system designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, inferring from the input it receives how to generate outputs such as predictions, content, recommendations, or decisions that influence physical or virtual environments”
AI Act	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)
Article 50	Article 50 of the AI Act, which lays down transparency obligations for certain AI systems
B2B	Business-to-business
Commission	European Commission
DSA	Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services (Digital Services Act)
Draft Guidelines	Draft Guidelines on the implementation of the transparency obligations for certain AI systems under Article 50 of Regulation (EU) 2024/1689, published by the Commission for stakeholder consultation
EU	European Union
GPAI	General-purpose AI
METI	Ministry of Economy, Trade and Industry, Japan
MeitY	Ministry of Electronics and Information Technology, Government of India
UCPD	Directive 2005/29/EC of the European Parliament and of the Council concerning unfair business-to-consumer commercial practices in the internal market (Unfair Commercial Practices Directive)

C. About the Cyril Shroff Centre for AI, Law & Regulation

The Cyril Shroff Centre for AI, Law & Regulation (the “CSCAILR”) is a leading hub and centre of excellence that strives to produce high-quality research and actionable policy recommendations. We engage directly with policymakers and connect academic scholarship with policy implementation with the objective of enabling India to develop state-of-the-art frameworks for the regulation and governance of AI.

CSCAILR is a pioneering initiative at Jindal Global Law School (JGLS), O.P. Jindal Global University (JGU), India's top-ranked and globally reputed law school. Built under a generous endowment by Mr. Cyril Shroff, one of India's leading lawyers, CSCAILR is uniquely positioned to partner with academics, policy makers, think-tanks, law firms, the judiciary and businesses, to produce insights tempered with multi-stakeholder viewpoints.

Although CSCAILR is based in India, the Draft Guidelines are of universal relevance for at least two reasons. First, Indian providers and deployers of AI systems may fall within the territorial scope of the EU AI Act where the output produced by those systems is used in the Union, and, where the relevant conditions in Article 50 are met, may be directly subject to its transparency obligations. Second, the technical and interpretive choices made in the Draft Guidelines may influence regulatory thinking in other jurisdictions, including India, through dynamics commonly described as the “Brussels effect”. It is in that spirit, and not as a substitute for the perspectives of European stakeholders, that we offer the following comments.

D. Comments Submitted to the European Commission

The comments set out in this Part are those that the Centre submitted to the European Commission through the AI Office's targeted stakeholder consultation on the Draft Guidelines, using the online contribution form provided for that purpose (Contribution ID 36433eca-f2b0-44c5-9887-1d84d8933c9c, submitted on 3 June 2026). The form posed specific questions, organised by the section of the Draft Guidelines to which they related, and confined each response to a fixed character limit. What appears below is the text of those responses as submitted, reproduced under the questions to which they respond. The fuller analysis from which these responses were drawn, developed without the constraints of the form, follows in Part E.

SECTION II – OVERVIEW OF TRANSPARENCY OBLIGATIONS AND HORIZONTAL TOPICS

Question A. Are there any aspects, sections or paragraphs that require further clarification or refinement? (Yes)

Two horizontal matters would benefit from refinement.

The comments set out in this Part are those that the Centre submitted to the European Commission through the AI Office's targeted stakeholder consultation on the Draft Guidelines, using the online contribution form provided for that purpose (Contribution ID 36433eca-f2b0-44c5-9887-1d84d8933c9c, submitted on 3 June 2026). The form posed specific questions, organised by the section of the Draft Guidelines to which they related, and confined each response to a fixed character limit. What appears below is the text of those responses as submitted, reproduced under the questions to which they respond. The fuller analysis from which these responses were drawn, developed without the constraints of the form, follows in Part E.

SECTION II – OVERVIEW OF TRANSPARENCY OBLIGATIONS AND HORIZONTAL TOPICS

Question A. Are there any aspects, sections or paragraphs that require further clarification or refinement? (Yes)

Two horizontal matters would benefit from refinement.

First, Article 50 does not apply to GPAI models, and the Draft Guidelines only “encourage” model-level marking. Yet as paragraph 78 acknowledges, no single technique satisfies all four criteria of effectiveness, robustness, reliability and interoperability, and effective marking often requires coordination between the model and system layers. A downstream system provider building on a model that lacks the necessary technical means of marking faces a materially harder task under Article 50(2), and may be confined to surface-form techniques most vulnerable to removal. This burden falls hardest on smaller providers and startups relying on third-party models, and risks favouring vertically integrated providers. We submit that at the very minimum, the gesture in paragraph 24 towards Article 55(1)(b) be developed into a concrete expectation that providers of GPAI models with systemic risk make the necessary technical means of marking available to downstream providers on reasonable, transparent and non-discriminatory terms. This requires no legislative amendment.

Second concerns extra-territorial scope. Paragraphs 10 and 11 correctly confirm that third-country providers and deployers fall within scope where output is used in the Union, so many Indian providers and deployers are directly affected. The practical compliance pathway is, however, under-elaborated. In particular, the Draft Guidelines do not address whether a third-country deployer already subject to an analogous home-jurisdiction regime, such as India's 2026 IT Amendment Rules on

synthetically generated content, may rely on that compliance, or whether parallel compliance is required. Clarity on this equivalence question would reduce compliance friction without weakening the Article 50 standard.

SECTION III – TRANSPARENCY FOR INTERACTIVE AI SYSTEMS (ARTICLE 50(1))

Question A. Are there aspects requiring further clarification for effective implementation and compliance? **(Yes)**

Paragraphs 39 to 42 (Section 3.2.1) anchor the Article 50(1) “obviousness” exception in the UCPD “average consumer” standard. The importation is elegant, but its suitability for the specific epistemic problem of AI interaction warrants reflection. The average-consumer test was developed to assess whether a reasonably attentive person would be misled by commercial information. Identifying whether an interlocutor is human or artificial is a different, perceptual task, and contemporary empirical research converges on a sobering finding that even digitally literate users often fail to identify advanced conversational AI, particularly where the system mirrors conversational register, exhibits hesitation or simulates affect.

This fragility is compounded by structure. On the face of paragraph 41, the obviousness assessment is conducted by the provider, which has an interest in concluding that obviousness is satisfied. EU consumer law has historically treated such situations with caution, including by allocating the burden of proof to the trader.

We respectfully submit three refinements. First, the assessment should be expressly evidence-based, with providers expected to support reliance on the exception through user studies or comparable material calibrated to the actual target audience. Second, the burden of establishing that the threshold is met should rest unambiguously on the provider. Third, where evidence is mixed or unavailable, the default should be that the disclosure obligation applies, so the exception functions as a real exception rather than a parallel route to non-disclosure. These refinements are consistent with the purpose stated in paragraph 25 of enabling natural persons to calibrate their trust in AI interactions.

Question B. Are there additional practical examples or use-cases that should be clarified? **(Yes)**

Yes. Two clusters of use-cases would benefit from clarification.

First, the real-time content carve-out in paragraph 82. The exception for “real-time content generation that is ephemeral and consumed at the moment without being stored or disseminated further” assumes that persons are aware the content is AI-generated. This imports a contextual-awareness test overlapping substantially with the Article 50(1) obviousness exception, and the same empirical fragility applies. The example of video games is no longer reliably low-stakes. Generative AI is increasingly present in serious games used for education, vocational training and therapeutic contexts, where the assumption of user awareness is weakest and the consequences of mistaken belief are highest. The Draft Guidelines would benefit from clarifying how the carve-out applies, if at all, to these higher-stakes environments, and from confirming that ephemerality alone does not displace disclosure where awareness cannot be assumed.

Second, advanced conversational systems calibrated for human-likeness. A use-case worth addressing is the system deliberately configured to mirror conversational register, insert hesitation markers or simulate emotional affect. These design choices directly reduce the likelihood that even a reasonably well-informed person will perceive AI-origin, and they are precisely the configurations on which empirical identification rates approach chance. Guidance confirming that more human-like configuration raises, rather than lowers, the disclosure expectation would give providers a clear and administrable signal.

SECTION VI – LABELLING OF DEEP FAKES AND CERTAIN TEXT PUBLICATIONS (ARTICLE 50(4))

Question A. Are there aspects requiring further clarification or refinement to ensure effective implementation? **(Yes)**

Two parts of Section 6 would benefit from refinement.

First, the attenuated regime for “evidently artistic, creative, satirical or fictional” deep fakes (Section 6.1.3, paragraphs 111 to 116). The contextual test for “evidently” relies on format, platform and audience expectation. The difficulty is that content often does not stay where it is first produced or shared. A satirical deep fake can be screenshotted, stripped of metadata, cropped to remove a watermark and recirculated where the original cues no longer reach the viewer, leaving a risk of deception identical to a non-satirical deep fake. We submit that the attenuated regime be conditioned on the disclosure being reasonably designed to survive foreseeable downstream re-sharing. Contextual cues alone should be insufficient where redistribution is foreseeable; an in-content disclosure that travels with the file (a visible watermark, an uncroppable on-screen label, or audio cueing) should be required, and the deployer-level obligation to maintain a perceptible disclosure should persist.

Second, the editorial responsibility exception for AI-generated text (Section 6.2.3, paragraphs 125 to 128). The current test is sound but assumes a recognisable publisher institution. We submit three clarifications. Where one natural person both authors and reviews, the exception should require evidence that review was a discrete step distinct from authoring. In institutions lacking a formal editorial hierarchy, the test should focus on whether a competent person has in substance exercised the functions in paragraphs 126 and 128, not on a named editor-in-chief. In hybrid workflows, the Draft Guidelines should specify how the person holding editorial responsibility is identified.

Question B. Are there additional examples or use-cases that should be clarified? (Yes)

Yes. Examples on both limbs of Article 50(4) would help.

A worked example should clarify the deployer's position where a satirical deep fake, lawfully published with full context on a comedy platform, is foreseeably re-shared on a closed messaging platform stripped of that context. Confirming that the deployer must build in a portable, perceptible disclosure in such scenarios would make the attenuated regime workable.

On editorial responsibility for AI-generated text, as noted above, three publication models increasingly common in practice fall outside the implicit “newspaper with an editor-in-chief” frame and would benefit from worked examples. First, the creator-economy newsletter sent from a personal Substack or comparable platform, where the author is also the de facto editor. Second, the institutional blog run by a small organisation such as an NGO, research centre or professional body without a formal editorial hierarchy. Third, the hybrid workflow in which AI-generated text passes through a mix of automated and human checks whose substantive depth is contested. Examples showing how the test applies to each would give these widespread, lower-formality publishers a clear compliance path and prevent the exception from turning on the presence of traditional newsroom structures.

E. Detailed Comments

Article 50 of the AI Act stands at the intersection of several distinct legal traditions: consumer protection law's concern with informational asymmetry, media law's interest in editorial accountability, copyright and personality-rights regimes' protection of identity, and the more recent body of platform regulation aimed at safeguarding the integrity of the public information ecosystem. The Draft Guidelines necessarily attempt to mediate among these traditions. Our detailed comments below proceed by identifying, in each instance, the existing interpretive framework offered by the Draft Guidelines, the gap or fragility we perceive, and a constructive expansion or refinement that we respectfully submit for the Commission's consideration.

1. THE “OBVIOUSNESS” EXCEPTION AND THE AVERAGE-PERSON STANDARD (SECTION 3.2.1)

- 1.1. Article 50(1) of the AI Act exempts providers from the disclosure obligation where it is “obvious from the point of view of a natural person who is reasonably well-informed, observant and circumspect” that they are interacting with an AI system. The Draft Guidelines, at paragraphs 39 to 42, explicitly anchor this standard in the established notion of the “average consumer” drawn from EU consumer-protection jurisprudence under the UCPD. The conceptual elegance of this importation is undeniable: the average-consumer test is a familiar, judicially well-developed standard, and its incorporation provides interpretive continuity across instruments. We respectfully submit, however, that the standard's suitability for the specific epistemic problem posed by advanced AI interaction warrants further reflection.
- 1.2. The average-consumer test was developed in a context where the question to be assessed by the hypothetical addressee was, in essence, a question about commercial information: would a reasonably attentive person be misled by this advertisement, label or claim? The cognitive task involved in identifying whether an interlocutor is human or artificial is qualitatively different. It is closer to a perceptual or attributional task than to a comprehension task, and it is one on which contemporary empirical research is converging towards a sobering conclusion: even digitally literate users routinely fail to identify advanced conversational AI, particularly where the system is configured to mirror conversational register, exhibit hesitation or simulate affect.¹ The concern is that, unless operationalised through evidence-based assessment, the obviousness exception may risk overestimating the capacity of even reasonably informed users to identify AI-origin in advanced conversational settings.
- 1.3. This empirical fragility is compounded by a structural feature of the exception. The obviousness assessment is, on the face of paragraph 41 of the Draft Guidelines, to be conducted by the provider. The provider has both an informational advantage (it knows precisely how human-like its system is configured to be) and an interest in concluding that obviousness is satisfied (since the conclusion exempts it from the disclosure obligation). This combination is one that EU consumer law has historically treated with caution, including by allocating the burden of proof on certain claims to the trader rather than the consumer.

¹ See, illustratively, Jeffrey T Hancock, Mor Naaman, Karen Levy, AI-Mediated Communication: Definition, Research Agenda, and Ethical Considerations, *Journal of Computer-Mediated Communication*, Volume 25, Issue 1, January 2020, Pages 89–100, <https://doi.org/10.1093/jcmc/zmz022>; and the more recent line of Turing-test-style empirical evaluations of GPT-class systems published since 2023, which consistently report identification rates close to chance for sustained interactions. See, for instance, Cameron Robert Jones, Ishika Rathi, Sydney Taylor, and Benjamin K. Bergen, 2025, People cannot distinguish GPT-4 from a human in a Turing test, in *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*, Association for Computing Machinery, New York, NY, USA, 1615–1639. <https://doi.org/10.1145/3715275.3732108>

- 1.4. The video-game carve-out in paragraph 82 also raises a similar concern. The exception provided to “Real-time content generation that is ephemeral and consumed at the moment without being stored or disseminated further” assumes that “the persons are aware that the content is AI-generated or manipulated”. It appears to import a contextual-awareness test that overlaps significantly with the Article 50(1) obviousness exception, and the same empirical fragilities apply. The increasing presence of generative AI in serious games used for education, training and even therapeutic contexts means that the gaming environment can no longer be assumed to be a low-stakes informational space.
- 1.5. It is submitted that the Draft Guidelines should be revised in three respects. First, the obviousness assessment should be made expressly evidence-based rather than purely conceptual: providers seeking to rely on the exception should be expected to support that reliance with appropriate user studies or other empirical material, calibrated to the actual target audience of the system. Second, the burden of establishing that the obviousness threshold is met should rest unambiguously on the provider. Third, where empirical evidence is mixed or unavailable, the default position should be that the disclosure obligation applies; the exception should function as an exception, not as a parallel route to non-disclosure. These refinements are consistent with the Draft Guidelines' own stated purpose, articulated in paragraph 25, of enabling natural persons to “calibrate their trust” in AI interactions.

2. ATTENUATED DISCLOSURE FOR “EVIDENTLY ARTISTIC, CREATIVE, SATIRICAL OR FICTIONAL” DEEP FAKES (SECTION 6.1.3)

- 2.1. Article 50(4), first subparagraph, of the AI Act provides for an attenuated transparency obligation in respect of deep fakes that form part of “evidently artistic, creative, satirical, fictional or analogous” works or programmes. The Draft Guidelines, at paragraphs 111 to 116, develop a contextual test for the “evidently” criterion, drawing on the format and style of the content, the platform or medium on which it is presented, and audience expectations. As an interpretive matter, this is consistent with how comparable contextual carve-outs operate in EU media law and copyright (for example, in respect of pastiche or parody).
- 2.2. The difficulty is that contemporary content does not stay where it is born. A satirical deep fake published on a comedy website with a clear stylistic and contextual frame can, within minutes, be screenshotted, stripped of its surrounding metadata, cropped to remove its watermark, and recirculated on a messaging platform where the original context is no longer perceivable. At that downstream stage, the very contextual cues that the Draft Guidelines treat as constitutive of “evidently artistic” character no longer reach the natural person exposed to the content. The risk of deception is then identical to that posed by a non-satirical deep fake, even though the original deployer remains within the attenuated regime.
- 2.3. This is not a hypothetical concern. The Indian experience with messaging-platform-mediated information disorder, particularly on WhatsApp during electoral and communal-tension periods, has been studied extensively and demonstrates the speed and severity with which decontextualised content can drive real-world harm.² The Draft Guidelines themselves recognise this dynamic obliquely in Recital 134 of the AI Act and in the interplay with Article 35(1)(k) of the DSA discussed at paragraph 118, but they do not yet draw a corresponding consequence at the level of the deployer's own labelling obligation.
- 2.4. It is submitted that the attenuated regime under Section 6.1.3 should be conditioned on the disclosure measure being reasonably designed to survive foreseeable downstream re-sharing. In practical terms, this would mean that contextual cues alone, such as the platform on which the deep fake is published or its embedding in a programme of evidently satirical character, are insufficient. The deployer should additionally be required to incorporate an in-content disclosure that travels with the file (for example, a visible watermark, an on-screen label that cannot be cropped away without destroying the content, or

² Timothy McLaughlin, 'How WhatsApp Fuels Fake News and Violence in India' Pulitzer Center (13 December 2018), <https://pulitzercenter.org/stories/how-whatsapp-fuels-fake-news-and-violence-india>

audio cueing) where the deep fake is reasonably likely to be redistributed on third-party platforms. This formulation preserves the legitimate space for satirical, fictional and creative work while preventing the attenuated regime from becoming, in effect, a route by which deep fakes enter the broader information ecosystem unmarked.

- 2.5. A complementary clarification would help. The Draft Guidelines should make explicit that where a deployer reasonably foresees that the deep fake will be redistributed beyond the original platform of publication, the obligation to maintain a perceptible disclosure persists at the deployer level, even where downstream redistributors may also be subject to their own obligations under the DSA or other instruments. The chain of responsibility should not be permitted to dilute at each handover.

3. THE EDITORIAL RESPONSIBILITY EXCEPTION FOR AI-GENERATED TEXT (SECTION 6.2.3)

- 3.1. Article 50(4), second subparagraph, exempts AI-generated text from the disclosure obligation where it has undergone human review or editorial control and where a natural or legal person holds editorial responsibility for its publication. The Draft Guidelines, at paragraphs 125 to 128, draw on the European Media Freedom Act's definition of editorial responsibility to flesh out the test, and they helpfully clarify that superficial, formal or procedural checks are insufficient. This anchoring is welcome.
- 3.2. The framework, however, assumes a recognisable publisher institution. It maps well onto a newspaper with an editor-in-chief, an academic publisher with peer review, or a public authority issuing communications under institutional sign-off. It maps less well onto three increasingly significant categories of publication: first, the creator-economy newsletter sent from a personal Substack or comparable platform, where the author is also the de facto editor; second, the institutional blog or website operated by a small organisation (an NGO, a research centre, a professional body) without a formal editorial hierarchy; and third, the hybrid publication in which AI-generated text passes through a workflow of automated and human checks the substantive depth of which is contested.
- 3.3. It is submitted that the Draft Guidelines could usefully provide additional clarification on three points. First, where a single natural person performs both the authoring and the review function, the exception should be available only where there is evidence that the review was a discrete cognitive step distinct from the authoring (for example, a time-separated re-reading with documented substantive changes). Second, in institutional contexts lacking a formal editorial hierarchy, the test should focus on whether a competent natural person has, in substance, exercised the functions described in paragraphs 126 and 128, rather than on the existence of a named editor-in-chief. Third, in hybrid workflows, the Draft Guidelines should similarly outline a mechanism for identification of the person holding the editorial responsibility.

4. INTERPLAY WITH THE GENERAL-PURPOSE AI MODEL LAYER (SECTION 2.6)

- 4.1. Section 2.6 of the Draft Guidelines correctly observes that Article 50 “does not explicitly apply to GPAI models” and confines itself to “encouraging” model-level marking by GPAI model providers in order to facilitate downstream compliance with Article 50(1) and (2). This drafting is faithful to the literal text of Article 50. We respectfully submit, however, that the regulatory architecture as currently expressed leaves an integration gap that materially affects the achievability of the Article 50 transparency objectives.

- 4.2. The technical reality of marking and detection, as the Draft Guidelines acknowledge at paragraph 78, is that no single technique presently satisfies all four quality criteria of effectiveness, robustness, reliability and interoperability. Effective marking architectures are therefore likely to require coordination between the model layer and the wider AI-system layer. Token-level watermarking may need to be implemented at, or very close to, the point of generation, while provenance metadata, cryptographic signatures, and lifecycle records are typically implemented through the surrounding system architecture that creates, stores, signs, and transmits information about the generated content. A downstream system provider that builds on a GPAI model lacking these primitives faces a significantly more difficult compliance task, and may be limited to techniques operating on the surface form of the output (such as visible watermarks on images), which are also the techniques most vulnerable to adversarial removal.
- 4.3. This asymmetry has important implications for both regulatory effectiveness and market structure. The providers best positioned to implement transparency-by-design measures at scale are often GPAI model providers, yet Article 50 places the operative compliance burden primarily on downstream system providers. This creates a misalignment between technical capability and regulatory responsibility. The burden is likely to fall most heavily on smaller providers, startups, and other actors that rely on third-party GPAI models and lack the ability to modify model-level architecture. As a result, compliance may become more costly and less effective, while inadvertently favouring vertically integrated providers that control both the model and system layers. Greater guidance on responsibility-sharing across the AI value chain would therefore improve both the effectiveness and proportionality of the Article 50 framework.
- 4.4. The Draft Guidelines do gesture, in the second sentence of paragraph 24, towards the possibility that model-level transparency measures may form part of the systemic-risk mitigation obligation under Article 55(1)(b) for GPAI models with systemic risk. It is submitted that this gesture should be developed into a more concrete expectation. Specifically, the Draft Guidelines could clarify that providers of GPAI models with systemic risk will be expected, as part of their Article 55(1)(b) obligations, to make available to downstream system providers the marking primitives necessary for those providers to comply with Article 50(2), and to do so on reasonable, transparent and non-discriminatory terms. Such a clarification would close the integration gap, would not require any legislative amendment, and would be consistent with the Commission's general approach to interoperability in the digital sphere.

5. EXTRA-TERRITORIAL APPLICABILITY AND THE BRUSSELS EFFECT

- 5.1. Article 2(1)(a) and (c) of the AI Act establish that providers and deployers located in third countries fall within the scope of the AI Act where the output of the AI system is used in the Union. The Draft Guidelines confirm this in paragraph 10 and again in paragraph 11. As a practical matter, this expansive jurisdictional approach means that a significant number of Indian developers, providers, deployers, and service providers may be subject to the transparency obligations under Article 50, notwithstanding the absence of any physical establishment within the European Union.
- 5.2. First, the territorial reach of Article 50 is, we respectfully submit, an under-acknowledged dimension of the Draft Guidelines. While the Draft Guidelines correctly state the legal position, the practical compliance pathway available to third-country providers and deployers is not elaborated. In particular, the Draft Guidelines do not address whether a third-country deployer that is also subject to an analogous (but not identical) disclosure regime in its home jurisdiction can rely on compliance with that home regime for the purposes of Article 50, or whether parallel compliance will be required. India, for instance, has begun developing its own approach to synthetically generated content under the 2026 IT Amendment Rules. Clarity on the equivalence question would materially reduce compliance friction without weakening the Article 50 standard.
- 5.3. Second, and consequentially, the technical solutions developed by EU-compliant providers for the marking and detection obligation under Article 50(2) are likely to become global defaults, irrespective of whether other jurisdictions formally require the same techniques. This is the classical Brussels-effect

dynamic, as developed in the work of Anu Bradford and others.³ The implication is that the technical standards endorsed by the AI Office in due course will, in effect, shape the global ecosystem for AI-generated content provenance. This is a position of significant responsibility, and we respectfully submit that it should be exercised with proportionate consultation.

- 5.4.** Third, it is submitted that the AI Office should, at this early stage and before finalising the technical standards contemplated in Section 4.2.3 of the Draft Guidelines, engage in interoperability dialogue with non-EU regulators. Constructive engagement with the Indian MeitY, the UK's Department for Science, Innovation and Technology, Singapore's Infocomm Media Development Authority, Japan's METI and AI Safety Institute, and counterpart bodies in jurisdictions developing their own AI-content provenance frameworks would have several benefits: it would avoid fragmentation of marking and detection ecosystems, it would reduce the compliance burden on global providers that would otherwise face divergent technical requirements, it would help legitimise the standards eventually adopted, and it would model a form of international regulatory cooperation appropriate to the cross-border character of AI-generated content.

³ A. Bradford, *The Brussels Effect: How the European Union Rules the World* (Oxford University Press, 2020).

F. Conclusion

Article 50 of the AI Act represents a significant and, in our respectful submission, broadly well-conceived attempt to translate a wide range of pre-existing regulatory traditions into a single horizontal transparency regime. The Draft Guidelines do important work in giving operational content to that regime. Our comments are accordingly offered in a spirit of constructive engagement. The recurring theme of our submissions is that, in several places, the Draft Guidelines lean on tests and constructs that may not bear the empirical or structural weight that the Article 50 objectives require: the average-person standard for obviousness, the contextual test for evidently artistic deep fakes, the institutional framing of editorial responsibility, and the soft treatment of the GPAI model layer. In each case we have suggested a constructive refinement that, in our view, would strengthen the Article 50 regime without disrupting its overall architecture.

We have also drawn the Commission's attention to the considerable extra-territorial reach of Article 50, particularly in relation to Indian providers and deployers, and to the Brussels-effect dynamic by which the technical standards endorsed under Article 50(2) will likely shape global practice. We respectfully submit that this dynamic counsels for early interoperability dialogue with non-EU regulators.

We thank the Commission for the opportunity to contribute to this consultation.



O.P. JINDAL GLOBAL
INSTITUTION OF EMINENCE DEEMED TO BE
UNIVERSITY
A Private University Promoting Public Service



CYRIL SHROFF
CENTRE FOR AI,
LAW & REGULATION
O.P. JINDAL GLOBAL UNIVERSITY



Jindal Global Law School
India's First Global Law School

 O.P. Jindal Global University, Sonipat Narela Road,
Sonipat-131001, (Delhi NCR), Haryana, India

 jindalglobaluni

 jindalglobaluni

 jindalglobaluni

 jindalglobaluni

 jguvideo

JGU - An Initiative of Jindal Steel & Power Foundation

 WWW.JGU.EDU.IN