Oñati Socio-Legal Series (ISSN: 2079-5971)
Oñati International Institute for the Sociology of Law
Avenida Universidad, 8 Apdo. 28
20560 Oñati, Gipuzkoa, Spain
Tel. (+34) 943 783064 / opo@iisj.net / https://opo.iisj.net
(cc) BY-NC-ND

From deliberative practice to algorithmic theatre: Generative AI and the risk of performative compliance in environmental due diligence

OÑATI SOCIO-LEGAL SERIES FORTHCOMING
DOI LINK: HTTPS://DOI.ORG/10.35295/OSLS.IISL.2836
RECEIVED 4 JULY 2026, ACCEPTED 3 NOVEMBER 2026, FIRST-ONLINE PUBLISHED 9 NOVEMBER 2026

IBRAR AHMAD[1]
AFTAB HAIDER[2]
ZHANG HUI[3]

Abstract

Generative AI is usually defended in the language of efficiency. Environmental due diligence, however, rests on documents machines can now draft, submit and summarise. This article proposes algorithmic theatre as an explanatory model of a risk in which compliance artefacts keep their legal form and authority while the deliberation they record thins or disappears. Using Dryzek's criteria of authenticity, inclusiveness and consequentiality, it reads three episodes as plausibility probes (AI pilots under the National Environmental Policy Act, fabricated public comments and false data in carbon credit verification) and specifies three hypothesised mechanisms: de-skilling, decoupling and legitimacy transfer. A comparison of the United States, the European Union and Pakistan suggests that exposure depends less on a jurisdiction's wealth than on whether its law attends to the provenance of compliance information, and is greatest where participatory infrastructure is already weak. Safeguards follow, including a draft provision for the disasters convention now under negotiation.

Key words

Algorithmic theatre, deliberative systems, environmental due diligence, generative AI, symbolic compliance, public participation, disaster law.

We thank the editor and both reviewers for their valuable feedback, which helped us improve the manuscript. The ideas, arguments, and analysis are our own. All authors have reviewed and take responsibility for the final version of the article.

[1] Ibrar Ahmad, Southwest University of Political Science and Law, Chongqing, China. Email: ibrarahmad557@gmail.com ORCID: https://orcid.org/0009-0001-0846-0218
[2] Aftab Haider, Southwest University of Political Science and Law, Chongqing, China. Corresponding Email: aftabhaider516@gmail.com ORCID: https://orcid.org/0009-0009-3339-5057
[3] Zhang Hui, Southwest University of Political Science and Law, Chongqing, China. Corresponding Email: cqzhanghui@swupl.edu.cn ORCID: https://orcid.org/0000-0001-7982-8780

This is an open access work licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license. (c) The Author(s)

Resumen

La IA generativa suele defenderse en términos de eficiencia. Sin embargo, la diligencia debida medioambiental se basa en documentos que las máquinas ya pueden redactar, presentar y resumir. Este artículo propone el «teatro algorítmico» como modelo explicativo de un riesgo en el que los artefactos de cumplimiento mantienen su forma jurídica y su autoridad, mientras que la deliberación que registran se diluye o desaparece. Utilizando los criterios de autenticidad, inclusividad y consecuencialidad de Dryzek, analiza tres episodios como pruebas de plausibilidad (proyectos piloto de IA en el marco de la Ley Nacional de Política Medioambiental, comentarios públicos falsificados y datos falsos en la verificación de créditos de carbono) y especifica tres mecanismos hipotéticos: la descalificación, el desacoplamiento y la transferencia de legitimidad. Una comparación entre Estados Unidos, la Unión Europea y Pakistán sugiere que la exposición depende menos de la riqueza de una jurisdicción que de si su legislación presta atención a la procedencia de la información sobre el cumplimiento, y es mayor allí donde la infraestructura participativa ya es débil. A continuación se presentan una serie de salvaguardias, incluida una propuesta de disposición para la convención sobre desastres que se está negociando actualmente.

Palabras clave

Teatro algorítmico, sistemas deliberativos, diligencia debida medioambiental, IA generativa, cumplimiento simbólico, participación ciudadana, derecho de catástrofes.

Table of contents

1. Introduction 4
2. Method, scope and definitions 5
    2.1. Four points of entry and three terms 5
    2.2. Plausibility probes: method and case selection 5
3. Due diligence and the premise it never states 6
    3.1. An obligation with four dimensions 6
    3.2. The deliberative dimension as an analytical grid 7
4. Three episodes read through the grid 8
    4.1. Assessment: AI in environmental review under NEPA 8
    4.2. Consultation: from forged comments to machine-read ones 9
    4.3. Verification: false data in carbon credit markets 10
5. Algorithmic theatre as an explanatory model 11
    5.1. What symbolic compliance already explains, and what it does not 11
    5.2. Front region, back region 11
    5.3. Three hypothesised mechanisms 12
    5.4. Conditions and indicators 13
6. Comparative exposure: The United States, the European Union and Pakistan 13
    6.1. The United States: documented incidents, a framework in flux 13
    6.2. The European Union: a comprehensive statute with partial reach 14
    6.3. Pakistan: thin participation and an early-stage AI policy 14
7. Beneficial uses, safeguards and the disasters convention 16
    7.1. Analysis and substitution 16
    7.2. Four safeguards 16
    7.3. A proposal for the convention on the protection of persons in the event of disasters 17
8. Two objections 18
    8.1. “This is fraud at a larger scale” 18
    8.2. “The deliberative framework is too idealised” 18
9. Conclusion 18
References 19
    Case law 24

> All the world's a stage / And all the men and women merely players.
> (William Shakespeare, *As You Like It*, II.vii)

1. Introduction

In 2017, the United States Federal Communications Commission (FCC) asked the public what it thought of a proposal to repeal the net neutrality rules, and more than 22 million comments arrived. The New York State Attorney General (OAG) later found that nearly 18 million of them, some 82 per cent, were fake (OAG 2021a). They came from opposite camps. Broadband companies paid commercial lead generators, which drew consumers in with gift cards and sweepstakes and then filed comments in their names without consent; that campaign alone produced more than 8.5 million fake comments. Separately, a single student used automated software to submit most of 9.3 million comments supporting net neutrality under invented identities (OAG 2021a). The proceeding met every formal marker of public participation. Much of what it recorded was counterfeit. The episode predates generative AI, and that is the reason to begin with it: it shows how a participation channel behaves once the cost of producing inputs falls, before large language models lowered that cost again.

Environmental due diligence runs through channels of the same kind. Impact assessments, consultation records and monitoring reports are documents, and generative AI now drafts, submits, summarises and checks documents of this sort. Federal agencies in the United States are piloting AI tools in environmental review (Saren 2025); governments have begun to let AI read consultation responses (Department for Science, Innovation and Technology —DSIT— 2025). That organisations can adopt the forms of law without its substance is old news to the literature on symbolic compliance (Edelman 1992, Krawiec 2003). The question here is narrower. Under what institutional conditions does the use of generative AI in producing and processing environmental due-diligence artefacts risk preserving the procedural form of compliance while weakening the deliberation from which that form draws its legitimacy, and how do existing legal frameworks respond?

The answer offered is an explanatory model rather than an empirical finding. The term proposed is *algorithmic theatre*: compliance artefacts that carry the authority of legal procedure but are generated or processed by systems that take no part in deliberation, so that the procedure may perform a deliberation that did not occur. The model is a hypothesis about a risk and its mechanisms, not a claim that environmental due diligence has already been transformed. Its contribution is threefold: it identifies what is analytically distinct about AI-mediated performance compared with cosmetic compliance; it specifies the mechanisms through which the risk would operate and the indicators by which it could be tested; and it maps where the law of three jurisdictions attends, and fails to attend, to the provenance of compliance information. The distinction that organises the argument is an old one. A legal procedure can possess the appearance of legitimacy (the filed document, the logged comment, the certified credit) without its substance, which is the deliberation those forms were built to house. For most of the history of environmental law the two travelled together, because producing the appearance required doing the work. Generative AI loosens that link.

Section 2 sets out method, cases and limits. Section 3 places the deliberative dimension within the wider doctrine of due diligence and builds the analytical grid, which Section 4 applies to three episodes. Section 5 develops the model against the literature on symbolic compliance, Section 6 compares the three jurisdictions, Section 7 turns to beneficial uses and safeguards, including a proposal for the disasters convention, and Section 8 answers two objections.

2. Method, scope and definitions

2.1. Four points of entry and three terms

By *AI-mediated due diligence* this article means the use of generative or analytical AI, by any actor in the due-diligence chain, to produce or process the documentary artefacts through which the obligation is discharged. Four points of entry matter. The first is *authoring*, where a project proponent, consultant or agency drafts an assessment with AI. The second is *participating*, where third parties use AI to produce consultation submissions. The third is *processing*, where an authority uses AI to summarise or classify what it receives. The fourth is *certifying*, where verifiers or registries rely on data that AI has generated or checked. Asked whose use of AI matters, the analysis therefore gives four answers rather than one, and the risk differs at each point.

Two further terms carry weight. *Legitimacy* is used in Suchman's sense: "a generalized perception or assumption that the actions of an entity are desirable, proper, or appropriate within some socially constructed system of norms, values, beliefs, and definitions" (Suchman 1995, 574). The concern is with what Schmidt (2013) calls throughput legitimacy, the quality of the governance process itself, as distinct from the input of participation and the output of results. *Fraud* means a knowing misrepresentation of material fact made to induce reliance. The definition matters because, as Section 8 argues, algorithmic theatre can arise without it.

2.2. Plausibility probes: method and case selection

The study is a qualitative, theory-building analysis of the kind George and Bennett (2005) call a plausibility probe, in which cases establish whether a proposed mechanism is credible enough to deserve systematic testing rather than how often it occurs. The corpus consists of primary legal and official documents: the two international advisory opinions on climate change and the earlier case law on due diligence; the International Law Commission's (ILC) draft articles of 2001 and 2016; the EU AI Act and EIA Directive; the National Environmental Policy Act (NEPA) as amended in 2023 and the 2025 instruments affecting its regulations; Pakistan's 1997 Act, 2000 Review Regulations and 2025 AI policy; the New York Attorney General's findings; the Administrative Conference of the United States recommendation on mass comments; the 2024 enforcement release of the Commodity Futures Trading Commission (CFTC); and government evaluations of AI consultation tools. They were read deductively against Dryzek's (2009) three criteria of deliberative capacity (authenticity, inclusiveness and consequentiality) and inductively for the mechanisms through which each criterion came under strain.

Cases were chosen purposively against three criteria. Each corresponds to a different station in the due-diligence chain codified by the ILC: assessment, consultation and verification. Each is documented by an official investigation, an enforcement action or published legal analysis, so that the facts do not rest on inference. And together they cover different points of AI entry. The comment-fabrication episode is neither environmental nor generative, and it is retained deliberately as a baseline. Notice-and-comment under the US Administrative Procedure Act is the same architecture through which public comment on federal environmental rules and reviews operates, and the episode shows how that architecture behaves when the cost of inputs falls. Experimental work on generated text (Weiss 2019, Kreps and Kriner 2024) and recent experience with AI-analysed consultations then show what generative tools add.

The three jurisdictions are contrasting configurations, not a representative sample. The United States offers the most documented incidents and the most advanced AI pilots in environmental review. The European Union has the most comprehensive AI statute in force and a harmonised EIA regime. Pakistan is a Global South system with a formal EIA law since 1997, an evaluative literature that records weak participation, a newly adopted national AI policy, and primary legal materials to which the authors have direct access. It lets the analysis ask how the risk presents where participatory infrastructure is thin.

The limits should be stated with the same directness. The article does not measure how far AI is used in environmental due diligence, and no public data yet allow it. The cases illustrate; they are not representative in any statistical sense, and only one of them involves generative AI in an environmental procedure. For Pakistan no recorded instance of AI use in EIA was found, so the analysis there concerns exposure rather than observed practice. The model is offered for testing, and Section 5.4 sets out the indicators by which it could be confirmed or falsified.

3. Due diligence and the premise it never states

3.1. An obligation with four dimensions

Due diligence in international environmental law is a bundle of obligations rather than a single one. Its preventive dimension is an obligation of conduct, to deploy adequate means against significant harm. Its organisational dimension is the duty to establish and maintain a national system of rules, institutions and enforcement. Its technical dimension is a standard of care calibrated to scientific and technological knowledge. Its procedural dimension consists of the steps of assessment, notification, consultation and monitoring (Krieger *et al.* 2020, Ollino 2022). Each has authority behind it. In *Pulp Mills* the International Court of Justice (ICJ) identified the organisational element, holding that due diligence "entails not only the adoption of appropriate rules and measures, but also a certain level of vigilance in their enforcement and the exercise of administrative control applicable to public and private operators" (ICJ 2010, para. 197), and it treated environmental impact assessment as a requirement of general international law where significant transboundary harm is risked (para. 204). The Seabed Disputes Chamber called due diligence "a variable concept", since measures once sufficient may become insufficient "in light, for instance, of new scientific or technological knowledge" (International Tribunal for the Law of the Sea —ITLOS— 2011, para. 117).

The procedural dimension is codified in the ILC's 2001 articles on prevention, which require prior authorisation (article 6), an assessment of risk that includes environmental impact (article 7), notification (article 8), consultation (article 9), exchange of information while the activity continues (article 12) and information to the public likely to be affected (article 13) (ILC 2001). Paine (2026) calls the result "organised proceduralism", a scaffold of mechanisms that generate and process the information states need if they are to discharge their obligations of prevention. The ICJ's 2025 advisory opinion ties the standard to knowledge. Scientific information, the Court held, "may provide the necessary evidence to assess the probability and seriousness of possible harm, informing the required standard of due diligence" (ICJ 2025, para. 283), and due diligence "requires States to take not only substantive measures, but also certain procedural steps" (para. 295). ITLOS adds that a state must "establish a national system, comprising legislation, administrative procedures and an enforcement mechanism" (ITLOS 2024, para. 235), and the ILA Study Group treats information, notification, cooperation, impact assessment and monitoring as corollaries of the general obligation (ILA 2016). Each of these steps leaves a document behind, and it is the document that the law inspects.

Two complications bear on how AI enters this field. The first is definitional. No uniform definition of environmental due diligence holds across the regimes that use the term: in international law it denotes a standard of state conduct, in corporate regulation a firm's process for identifying and mitigating adverse impacts, and in transactional practice a buyer's investigation of environmental liabilities. AI tools are marketed across all three, and without a shared definition it is harder to say which safeguards apply to which artefact. The second complication is that environmental information can be drawn into the register of security. Securitisation theory describes how an issue is moved out of ordinary politics and presented as an existential threat that licenses exceptional measures (Buzan *et al.* 1998). The ILC anticipated one form of the resulting tension when it allowed a state of origin to withhold data "vital to the national security" while cooperating in good faith (ILC 2001, art. 14). Where environmental data are treated as security-sensitive, those who gather them may face legal exposure, independent verification becomes harder, and the gap between a compliance document and the conditions it reports is harder to inspect. Both complications widen the space in which an artefact can stand in for the practice it reports.

3.2. The deliberative dimension as an analytical grid

The procedural dimension is where the deliberative character of due diligence shows most plainly, and it is the dimension on which this article concentrates. The claim is not that due diligence is only deliberative. Its preventive, organisational and technical dimensions remain, and AI may affect each of them differently. The claim is rather that the procedural steps rest on a deliberative premise which the doctrine leaves implicit. Dryzek and Niemeyer developed the idea of discursive representation to reach a problem that electoral theory cannot: how governance can be democratically legitimate where there is no settled demos, no bounded people whose votes confer authority. Their answer is that democracy can represent discourses as well as persons (Dryzek and Niemeyer 2008). A legitimate system makes sure that the relevant discourses (scientific, economic, ecological and community-based) are engaged rather than excluded, and it does so across linked sites rather than in a single forum (Mansbridge *et al.* 2012).

Discursive representation, they add, is "particularly relevant to contexts like the international system, where there is no well-defined demos" (Dryzek and Niemeyer 2008, 481).

Dryzek's (2009) three criteria of deliberative capacity make the framework usable. Authenticity requires that deliberation induce reflection without coercion and connect particular claims to more general principles. Inclusiveness requires that the range of interests and discourses present in a setting be represented. Consequentiality requires that deliberation affect collective decisions. Each procedural step of due diligence can be read against these criteria, and participation in EIA has long been justified on exactly such normative and substantive grounds (Glucker *et al.* 2013). An impact assessment is authentic only if it reports an assessment that someone has reasoned through. A consultation is inclusive only if its submissions represent discourses that people hold. The sequence is consequential only if decisions respond to what assessment and consultation disclose.

The architecture rests on a premise the doctrine has never stated, because it has never had to. The premise is authenticity of inputs. Legitimacy flows through the deliberative chain on the assumption that the information entering it is genuine. Environmental law does not produce the knowledge on which it acts. It borrows that knowledge from the scientist's fieldwork, the community's testimony and the verifier's inspection, and its procedures discipline the borrowing. While those practices were costly to counterfeit, the dependence was tolerable. Generative AI lowers the cost both of counterfeiting the practices and of substituting for them. The law still borrows; it can no longer assume that the lender did the work.

4. Three episodes read through the grid

Each episode below is read against the three criteria, and each closes with what the evidence can and cannot show.

4.1. Assessment: AI in environmental review under NEPA

The clearest evidence that AI is entering environmental assessment comes from the United States. Saren (2025) documents federal pilots, led by the Department of Energy and the Pacific Northwest National Laboratory, that apply AI to NEPA documents, from searching past reviews to assisting with drafting and public engagement, and asks whether AI-generated environmental impact statements "can withstand judicial review". The legal setting has moved quickly. Since 2023 the statute itself has required agencies to "ensure the professional integrity, including scientific integrity, of the discussion and analysis in an environmental document" (42 U.S.C. §4332(2)(D)). In 2025 the Council on Environmental Quality removed its implementing regulations (CEQ 2025a), and the Supreme Court told lower courts to afford agencies substantial deference on the scope and depth of their reviews (*Seven County Infrastructure Coalition v. Eagle County* 2025).[4] The integrity duty survives in the statute, while the intensity of review that once enforced it has been reduced.

[4] The holding concerned the upstream and downstream effects an agency must consider and says nothing about AI. Its bearing here is on the intensity of review that any AI-assisted statement would now meet.

Read against the grid, the case puts authenticity under strain. Deep-learning systems can operate as "black boxes" that produce results without a clear account of how they were reached (Pasquale 2015, Burrell 2016, Saren 2025). A regime that demands demonstrable scientific integrity is thus asked to rely on a technology that does not give its reasons. Saren draws the doctrinal consequence: an agency unable to explain the rationale behind a model's conclusions risks failing the standard of reasoned decision-making, so that the quality of the data becomes a condition of the decision surviving review. Where that condition is not met, the regulator risks scrutinising a procedural artefact rather than a scientific judgment, and the public receives a document whose provenance it cannot see. The evidence shows adoption and a documented legal concern; it does not show that any statement has been accepted without the judgment it reports. Saren's own proposal, explainable systems with audit logs and expert review at critical stages, indicates that the risk is recognised and, in principle, manageable.

4.2. Consultation: from forged comments to machine-read ones

The FCC proceeding of 2017 remains the most fully documented instance of public participation manufactured at scale (Figure 1). What matters here is who produced the submissions and how. The industry-funded lead generators co-opted real identities harvested through incentive offers; the student invented identities with software (OAG 2021a). Neither campaign needed generative AI. Both needed money and coordination, and the broadband industry spent US$8.2 million on its own (OAG 2021b).

FIGURE 1

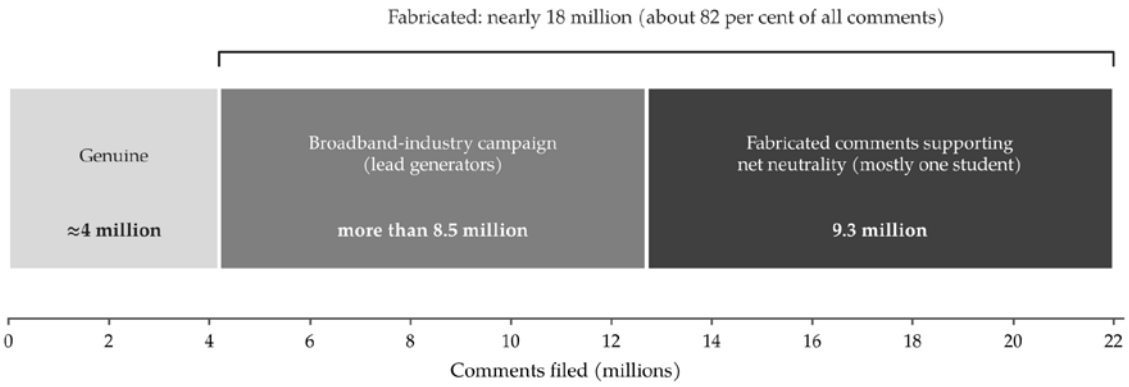


Figure 1. Composition of the more than 22 million comments filed in the FCC net neutrality proceeding, 2017.

Note: Shares are approximate. Most of the comments supporting net neutrality were submitted by one student using automated software. Source: OAG (2021a, 2021b).

Generative AI removes that constraint. In 2019 a student used a language model to submit 1,001 synthetic comments to a federal Medicaid waiver proceeding, 55.3 per cent of all comments received, and surveyed readers identified them at no better than chance (Weiss 2019). A field experiment found that state legislators answered AI-written constituent letters at rates only slightly below those for human-written ones (Kreps and Kriner 2024). Haider *et al.* (2024) found that 57 per cent of the questionable GPT-fabricated papers they retrieved from Google Scholar dealt with policy-relevant subjects,

among them environment and health, the fields on which environmental due diligence draws.

The deliberative worth of participation turns on authenticity rather than volume. In Dryzek and Niemeyer's terms, fabricated submissions are counterfeit discursive representation: they place in the record discourses that no participant holds, or inflate one discourse until the others are crowded out, and in either case they put inclusiveness under strain. The institutional response confirms that this is a problem of design and not only of wrongdoing. The Commission went ahead with its rulemaking. The Administrative Conference of the United States (ACUS) later recommended that agencies develop means of identifying and managing mass, computer-generated and falsely attributed comments (ACUS 2021). Agencies tend to respond to the substance of comments rather than their number (Balla *et al.* 2022), which limits the risk without removing it. A thousand copies of one argument remain one argument. The difficulty begins when a machine can write a thousand different ones.

Recent experience with AI-assisted consultation analysis in the United Kingdom and representational auditing research in Canada suggests that the challenge is no longer limited to fabricated submissions, but extends to AI-mediated interpretation of authentic public participation, raising new questions about deliberative legitimacy. In 2025 the UK government's Consult tool themed more than 2,000 responses to a Scottish Government consultation with results close to those of human reviewers, and the government announced that it intended to use the tool on major consultations without individual review of every response (DSIT 2025). An audit of the official summaries of Canada's 2025 national consultation on artificial intelligence found that, although average coverage did not fall below random baselines, clusters expressing dissent or rejection of AI were excluded at rates of 33 to 88 per cent (Mahajan 2026). Here the inputs are genuine and the risk sits at the point of processing, since the summary on which a decision rests may represent fewer discourses than the consultation contained. The average looked acceptable; much of the dissent did not survive the summary.

4.3. Verification: false data in carbon credit markets

The third episode moves from assessment and consultation to verification, the closing link in the due-diligence chain. In 2024 the Commodity Futures Trading Commission (CFTC) brought its first fraud action in the voluntary carbon credit market, charging the former chief executive of a project developer with supplying false information about the performance of emissions-reduction projects (CFTC 2024). A credit is the product of a chain of claims: a developer measures a reduction, an independent verifier reviews the measurement, and a registry issues the credit on the strength of the verification. Where measurement, verification and certification are all documentary acts, the whole chain of value can be corrupted at the level of information.

The case did not involve generative AI, and nothing here suggests that it did. Its relevance is structural, and the evidence on AI points in both directions. Tian and Shi (2025) find that corporate adoption of AI is associated with less greenwashing, partly through green innovation, which is a reminder that analytical AI can strengthen verification. The risk the model identifies arises elsewhere. Juvenal's question of who will guard the guards themselves (*Satires* VI) takes a new form once the guard and the

guarded can draft through the same class of system. Where one generative tool can produce both a performance report and the text that reviews it, the relationship between developer, verifier and registry loses the independence on which consequentiality depends, because the decision to issue a credit then responds to a record whose independence can no longer be assumed. Crawford (2021) adds a consideration that compliance accounting does not yet capture: the computation used to produce environmental claims carries energy, mineral and water costs of its own.

5. Algorithmic theatre as an explanatory model

5.1. What symbolic compliance already explains, and what it does not

The idea that compliance can be performed rather than practised is not new. Meyer and Rowan (1977) showed that organisations adopt formal structures as "myth and ceremony", decoupling them from technical activity to secure legitimacy. Edelman (1992) found that organisations respond to ambiguous law by building symbolic structures that signal compliance while leaving practice largely intact. Krawiec (2003) described cosmetic compliance, meaning internal structures that satisfy regulators and courts while doing little to reduce misconduct, and showed that such structures persist because legal decision-makers cannot readily tell effective programmes from cosmetic ones. Power (1997) traced the spread of rituals of verification in which audit produces comfort rather than knowledge. Any account of AI-mediated compliance must say what it adds to all this.

Three differences can be stated, each a matter of degree that may become a matter of kind at scale. The first concerns the performer. Symbolic compliance is produced by organisations with interests, which decouple form from practice by design or by drift. Algorithmic theatre can arise without any actor intending decoupling: a consultant who drafts an assessment with a language model may believe the result sound, and the system that wrote it holds no belief at all. The second concerns cost. Cosmetic compliance still has to be produced, and its cost limits its scale; generative AI reduces the marginal cost of a plausible artefact to almost nothing, as the contrast between the 2017 campaign and the 2019 Medicaid experiment shows. The third concerns the audience. In the symbolic-compliance literature the regulator who accepts the ceremony is human. Where authorities also use AI to process what they receive, a generated submission may be condensed by a generated summary, and no one in the chain need read the record whole. Symbolic-compliance theory explains why form and substance separate. The model adds an account of how the separation can occur without an author and be received without a reader. It also differs from algorithmic governmentality (Rouvroy and Berns 2013), which concerns how algorithms govern populations rather than the artefacts through which the governed are meant to take part.

5.2. Front region, back region

Goffman's dramaturgy explains why the theatrical figure is analytical rather than ornamental. Goffman distinguishes the sincere performer, who believes in the impression the performance fosters, from the cynical performer, who does not (Goffman 1959, 17-18). He separates the front region, where a performance is given, from the back region, where it is prepared and where the impression it fosters may be knowingly

contradicted (Goffman 1959, 112). Compliance documents are front-region performances addressed to an audience of regulators, courts and the public. Hilgartner (2000) showed that expert advisory bodies manage exactly this division, staging scientific advice for public audiences while controlling how much of its preparation can be seen. Generative AI enters the back region. It changes how the performance is prepared and leaves the front region as it was, so the audience sees no change.

The AI-generated artefact also occupies a position that Goffman's pairing did not anticipate. It is neither sincere nor cynical, because text produced by a language model is not grounded in communicative intent (Bender *et al.* 2021). Goffman observed, too, that audiences tactfully help performers sustain an impression, and an institution that treats a filed document as evidence of the assessment it reports extends that tact by default. Theatre is also a better figure than simulation. Baudrillard's (1994) simulacrum is a copy that has lost its original, whereas the environment that compliance documents describe is real; what may go missing is the act of judgment through which it was brought into the institution. Theatre keeps an audience, and that audience of regulators, courts and the public can still withdraw the legitimacy it confers.

5.3. Three hypothesised mechanisms

De-skilling names the displacement of professional expertise by algorithmic production, and it bears on authenticity. An impact statement has traditionally been written by scientists who gather field data, weigh impacts and exercise judgment. Where the statement is substantially machine-written, the relationship between science and regulation may change in kind and not only in degree: the regulator no longer assesses a scientific assessment but checks a procedural artefact whose scientific origins cannot be seen. The risk compounds over time. Expertise that is not exercised is not maintained, and the reviewer able to recognise an implausible assessment is a product of the practice that automation replaces.

Decoupling names the separation of the compliance artefact from the practice it claims to represent. It bears on inclusiveness and is the mechanism closest to Meyer and Rowan's. The FCC case shows it in its purest form, with the count of comments detached from the quality of engagement; the Canadian audit shows its processing variant, with a summary detached from the submissions it reports. Wardle and Derakhshan (2017) built their taxonomy of information disorder (misinformation, disinformation and malinformation) on falsity and intent to harm. AI-mediated compliance artefacts may fall outside all three categories. They can be procedurally correct, need not be false and may be produced with no intent to harm, and yet fail to report the practice they claim to report.

Legitimacy transfer bears on consequentiality, and it is the mechanism that most distinguishes the model from ordinary fraud. Fraud announces its own wrongfulness: there is a false statement, an identifiable wrongdoer and a prosecution to follow. In algorithmic theatre the impact statement is filed, the comments are logged and the credits are certified, and each performance takes place inside the institutional architecture of due diligence. Because the architecture still recognises the performance as one of its own, its legitimacy passes to the performance whether or not the deliberation the architecture was built to secure took place. Decisions remain formally

consequential, but they respond to the record rather than to the deliberation the record was meant to contain.

The three mechanisms are hypothesised to reinforce one another. De-skilling weakens the expertise that might notice a defective artefact; decoupling weakens the link that might expose it; legitimacy transfer ensures that an artefact nobody notices is treated as valid. If this interaction holds, the risk would tend to conceal itself, and its effects would accumulate gradually and out of view, much as Nixon (2011) describes for attritional environmental harm. Whether it holds is an empirical question.

5.4. Conditions and indicators

The model predicts that the risk is greatest where three conditions coincide: a high volume of documentary input relative to the capacity to review it; weak or late participation, which removes an independent check; and law that attends to the content of compliance artefacts but not to their provenance. Four kinds of evidence would support it: rising textual homogeneity across assessments prepared by different consultants for different sites; consultation records whose volume grows without a matching growth in distinct arguments; agency responses that address fewer site-specific claims; and verification reports that reproduce developer text. It would be weakened by evidence that AI-assisted assessments contain more site-specific data, and receive more searching review, than their predecessors. Each indicator can be measured with existing methods of text analysis. A model that cannot say what would refute it is a mood rather than a hypothesis.

6. Comparative exposure: The United States, the European Union and Pakistan

The technological pressure is much the same across jurisdictions. The legal response is not. The three systems are compared here on a single question: does the law attend to the provenance of the information on which environmental due diligence relies?

6.1. The United States: documented incidents, a framework in flux

The United States combines documented incidents with a framework in flux. It has recorded evidence both of AI entering environmental review (Saren 2025) and of public participation fabricated at scale (OAG 2021a). The statutory integrity duty addresses the content of environmental documents rather than the means of their production, and since 2025 there have been no government-wide NEPA regulations to elaborate it. Federal policy now promotes AI in permitting, and the Council on Environmental Quality's (CEQ) 2025 action plan encourages automated tools in environmental review (CEQ 2025b). Promotion is not itself the difficulty, since the analytical uses discussed in Section 7 may improve reviews. The difficulty is one of sequence: adoption is accelerating while the intensity of judicial review has been reduced, and no federal rule requires disclosure of AI authorship. In the FCC proceeding the scale of counterfeit participation came to light only through a later law-enforcement inquiry (OAG 2021b). The ACUS (2021) recommendation is the most developed response, and it concerns the processing of comments rather than environmental assessment.

6.2. The European Union: a comprehensive statute with partial reach

The European Union has a sophisticated framework whose reach into environmental compliance is partial. The AI Act (Regulation (EU) 2024/1689) is the most complete AI statute in force, yet AI used to draft environmental compliance documents is not among the high-risk uses listed in Annex III. Two provisions reach it indirectly. Providers of generative systems must mark synthetic outputs in a machine-readable format (Article 50(2)). Deployers who publish AI-generated text to inform the public on matters of public interest must disclose that fact, unless the text has undergone human review or editorial control and a person holds editorial responsibility (Article 50(4)).[5] A consultant who reviews and signs an AI-drafted environmental report would fall within that exception. The exception makes sense for publishing, but it leaves the regulator with no information about how the assessment was produced. The Act classifies AI by use and by risk to persons, and environmental compliance documentation escapes it because its effect is mediated: it passes through a regulatory decision before it reaches communities.

The EIA Directive offers a more promising hook. It requires the developer to ensure that the assessment report is prepared by "competent experts", and the competent authority to ensure that it has, or has access to, "sufficient expertise to examine" the report (Directive 2011/92/EU, as amended, art. 5(3)). Read purposively, a requirement of competence presupposes that the authority can know who exercised the expertise, and it could support a duty to disclose material AI assistance. The question grows as the Union digitalises compliance. The Corporate Sustainability Reporting Directive (Directive (EU) 2022/2464) mandates digital sustainability reporting; the Ecodesign Regulation (Regulation (EU) 2024/1781) introduces digital product passports; the Deforestation Regulation (Regulation (EU) 2023/1115) requires geolocated due-diligence statements filed through an information system. None of these instruments speaks to AI-generated content, although each builds a channel through which it can flow.

6.3. Pakistan: thin participation and an early-stage AI policy

Pakistan's inclusion is not an indictment. Similar patterns appear in many systems, and the purpose, as Section 2.2 explains, is to ask how the risk presents where independent checks are thin. The Pakistan Environmental Protection Act (PEPA) 1997 established a two-stage process: an Initial Environmental Examination (IEE) and, where significant adverse effects are likely, a full EIA, reviewed by environmental protection agencies that since the constitutional devolution of 2010 have operated largely at provincial level (PEPA 1997, s. 12). The 2000 Review Regulations require public notice and a public hearing for an EIA but not for an IEE (Pakistan EPA 2000). Evaluations of practice report that participation is largely confined to the statutory hearing, comes late and has limited influence on decisions, while review capacity is constrained (Nadeem and Hameed 2008, Nadeem and Fischer 2011). Arif *et al.* (2026) set these weaknesses in a wider frame, attributing the gap between Pakistan's climate policy architecture and its capacity to protect vulnerable populations to "fragmented governance, weak enforcement, and the absence of a rights-based approach".

[5] Whether a consultant who submits a report to an authority "publishes" it "to inform the public" at all is itself open to doubt, which narrows the provision's reach further.

On AI, the federal cabinet approved a National Artificial Intelligence Policy in July 2025 (Ministry of Information Technology and Telecommunication —MoITT— 2025). It is a policy rather than a statute, oriented to adoption, skills and public-sector use, and it does not address AI in environmental assessment. No recorded instance of AI use in Pakistani EIA practice was found, so the analysis concerns exposure rather than observed theatre. That exposure follows from the conditions set out in Section 5.4. An AI-drafted assessment submitted to an agency with limited review capacity, for a project whose public hearing comes late and changes little, meets fewer independent checks than it would in a system with early participation and specialised review. Review capacity is also a stock of skill. The Policy's commitment to train public servants in AI basics (MoITT 2025, item 2.2) addresses familiarity with the tool. The capacity to judge the scientific content of an assessment is human capital of another kind, and in the growth model of Berde *et al.* (2025) such capital accumulates through years of investment in education and depreciates when it is not renewed. Pakistan does not lack law. Like the other two jurisdictions, its law attends to the existence of an assessment more than to its provenance.

Set side by side, the three systems yield a lesson that Table 1 summarises. Vulnerability to algorithmic theatre does not track a jurisdiction's wealth or the sophistication of its statutes. What the three share is law that attends to the content of decisions, or to the technology behind them, more than to the provenance of the information on which decisions rest.

TABLE 1

Dimension	**United States**	**European Union**	**Pakistan**
EIA framework	NEPA 1969, amended 2023; CEQ regulations removed 2025; agency procedures.	EIA Directive 2011/92/EU (amended 2014); competent-expert requirement.	PEPA 1997; 2000 Review Regulations; provincial regimes since 2010.
Public participation	Notice-and-comment; mass fabrication documented (2017).	Aarhus Convention; structured, no provenance checks.	Public hearing for EIA only; late, limited influence.
AI regulation	No federal statute; 2025 policy promotes AI in permitting.	AI Act 2024; Article 50 transparency; environmental compliance not high-risk.	National AI Policy 2025 (non-statutory).
Recorded AI-related practice	AI pilots in NEPA review; synthetic comments (2019).	None specific to EIA identified.	None identified.

Dimension	United States	European Union	Pakistan
Principal exposure	Adoption outpacing review; reduced judicial scrutiny.	Mediated effects outside high-risk categories.	Thin participation and review capacity.
Existing legal hooks	NEPA §102(2)(D) integrity duty; ACUS Recommendation 2021-1.	EIA Directive art. 5(3); AI Act art. 50.	EPA review; statutory public hearing for EIA.
Overall exposure (qualitative)	High.	Moderate.	High, structural.

Table 1. Comparative exposure to algorithmic theatre in three jurisdictions.
(Note: Exposure ratings are qualitative judgments drawn from the analysis in Sections 4 to 6, not measurements.)

7. Beneficial uses, safeguards and the disasters convention

7.1. Analysis and substitution

The argument so far concerns a risk, and a fair assessment must set it against the benefits. AI already improves parts of environmental due diligence. Remote sensing combined with machine learning detects land-use change and emission sources at scales no inspector could match. Retrieval systems of the kind piloted under NEPA let reviewers search decades of prior assessments. Deduplication and clustering tools help agencies identify mass and computer-generated comments (ACUS 2021), and firm-level evidence suggests that adopting AI can reduce greenwashing (Tian and Shi 2025). AI can also help detect the misinformation it helps produce (Galaz *et al.* 2023). The relevant line runs between analytical uses, which interrogate data about the world, and generative uses that produce artefacts standing in for engagement with it. In the vocabulary of deliberative systems, the first can strengthen authenticity by giving deliberation better evidence, and the second can weaken it by substituting for deliberation. The line is not always sharp, since a generative tool can draft a section from verified data, and the safeguards below are designed to keep it visible rather than to prohibit a class of technology. The test is not what the tool is but whether its use can be seen.

7.2. Four safeguards

The literatures on algorithmic accountability and explainable AI supply the elements. Citron (2008) argued that automated administration requires technological due process: audit trails, notice and an opportunity to contest. Coglianese and Lehr (2017) showed that machine learning can be reconciled with administrative law if agencies can validate and explain their systems, and Kroll *et al.* (2017) showed how procedural regularity can be demonstrated by technical means. Explanation need not mean opening the model, since counterfactual explanations can show what would have changed an outcome

(Wachter *et al.* 2018). Accountability, in Bovens's (2007) sense, requires a forum before which an actor must explain and justify conduct, and internal audit frameworks spanning the life cycle of a system offer a practical way to supply one (Raji *et al.* 2020).

Translated into environmental due diligence, these suggest four safeguards, each tied to a criterion. For authenticity, disclosure of material AI assistance in any assessment, with an attestation by a named professional of the judgment exercised. For inclusiveness, provenance checks on consultation submissions proportionate to risk, and representational audits of any AI-generated summary of consultation responses, so that excluded discourses are visible to decision-makers (Mahajan 2026). For consequentiality, a record, in the decision itself, of how the authority engaged with site-specific evidence and with dissenting submissions. Across all three, the retention of logs sufficient to permit independent verification, in the spirit of the human oversight that Article 14 of the AI Act requires of high-risk systems. None of these needs new institutions. Each can be built into existing EIA review and comment procedures.

7.3. A proposal for the convention on the protection of persons in the event of disasters

International disaster law offers a timely opportunity to fix such safeguards in treaty form. In 2016 the ILC adopted draft articles on the protection of persons in the event of disasters. Article 9 requires each state to reduce the risk of disasters "by taking appropriate measures, including through legislation and regulations, to prevent, mitigate, and prepare for disasters", and it specifies that such measures include "the conduct of risk assessments, the collection and dissemination of risk and past loss information, and the installation and operation of early warning systems" (ILC 2016). The obligation is framed, like due diligence, as a duty to take appropriate measures, and its instruments are the same documentary artefacts examined here. In 2024 the General Assembly decided to elaborate a legally binding instrument on the basis of the draft articles by the end of 2027 (UNGA 2024); a preparatory committee met in 2026, and a diplomatic conference is scheduled to open in Bangkok in January 2027 (United Nations Office of Legal Affairs —UN OLA— 2026). As states rely more on AI for hazard modelling, risk mapping and early warning, the integrity of disaster-risk information becomes a condition of the protection the convention promises. The following text is offered for discussion, as an additional paragraph to the provision on disaster risk reduction:

> 3. In taking the measures referred to in paragraphs 1 and 2, each State shall ensure, as appropriate, that: (a) risk assessments and risk information are attributable to identifiable persons or institutions responsible for the professional judgment they contain; (b) the use of automated or generative systems in producing such assessments and information, or in analysing the views of persons likely to be affected, is disclosed; (c) the outputs of such systems are subject to meaningful human review before they inform decisions; and (d) records are kept that permit independent verification of the information relied upon.

The wording follows the draft articles' own preference for appropriate measures and leaves states latitude. It addresses the points of entry most relevant to disaster governance, namely authoring, processing and certifying, and it would give the convention a provision that reaches the information on which prevention depends without requiring states to regulate AI generally.

8. Two objections

8.1. "This is fraud at a larger scale"

The first objection reduces algorithmic theatre to cheaper fraud. It deserves its strongest form. Fraud law targets knowing misrepresentations by identifiable actors, and its remedy is to find and sanction the wrongdoer; some of the cases examined here fit that description, as the 2017 comments and the 2024 carbon credit action did. But the model also covers cases in which no one need intend to deceive: an assessment drafted with AI and signed in good faith, a summary that omits dissent because of how it was generated, a verification whose independence is compromised by shared tools. There is no fraud to prosecute in these cases, and the question Beckers and Teubner (2023) put, who answers when algorithms act within socio-digital institutions, becomes a question of institutional design. The sociology of law has long distinguished deviance, the individual breach of a norm, from institutional transformation, the change in the social character of the norms themselves. Algorithmic theatre, if the model holds, belongs to the second category, which is why prosecution alone would not reach it.

8.2. "The deliberative framework is too idealised"

The second objection charges the deliberative-systems framework with idealism: genuine deliberation, it says, has never existed in environmental governance, and the model merely restates a familiar gap between ideal and practice. The objection is partly right. Deliberative-systems theory does not require deliberation to be perfect. It requires the system to keep its capacity for deliberation, the procedures through which rival discourses can meet, however imperfectly. Even where participation is weak, a late public hearing still offers a forum in which community discourse can reach regulatory discourse. If AI drafts the assessment and summarises the responses without disclosure or review, that capacity is reduced further, and the reduction is harder to see. The model therefore diagnoses a possible loss of deliberative capacity rather than a failure to reach an ideal. The distinction has practical consequences. Strengthening participation and securing the provenance of the information on which it relies are different tasks, and both are needed.

9. Conclusion

This article has proposed algorithmic theatre as an explanatory model of a risk: that generative AI, entering environmental due diligence at the points of authoring, participating, processing and certifying, may preserve the procedural form of compliance while weakening the deliberation from which that form draws its legitimacy. Read through Dryzek's criteria, three documented episodes, in assessment, consultation and verification, make the model plausible without proving it, and the British and Canadian experience shows the risk moving from fabricated submissions to the interpretation of genuine ones. What separates the model from accounts of symbolic compliance is that it allows decoupling without an intending author and reception without a reader.

What cannot yet be said is how widespread the risk is. That question is empirical, and Section 5.4 sets out the indicators by which it can be answered. What can be said is that

the law of all three jurisdictions attends more to the content of compliance artefacts than to their provenance, and that the gap matters most where participation and review are already thin. The remedies are modest and available: disclosure, attestation, representational audit, recorded engagement and verifiable logs, and, in international law, a provision on the integrity of disaster-risk information in the convention now under negotiation. The FCC proceeding with which the article began was exposed because an investigator went behind the stage. The purpose of these safeguards is to make that inspection routine rather than exceptional.

References

Administrative Conference of the United States (ACUS), 2021. Recommendation 2021-1: Managing mass, computer-generated, and falsely attributed comments. *Federal Register* [online], 86, 36075. Available at: https://www.govinfo.gov/content/pkg/FR-2021-07-08/pdf/2021-14597.pdf

Arif, M. I. U., Mahsud, M. I., and Imran, S., 2026. From environmental governance to climate citizenship: Rethinking state authority, climate vulnerability and conflict dynamics in Pakistan. *Oñati Socio-Legal Series* [online], *16*(3), 1394-1419. Available at: https://doi.org/10.35295/osls.iisl.2556

Balla, S. J., *et al*., 2022. Lost in the flood?: Agency responsiveness to mass comment campaigns in administrative rulemaking. *Regulation & Governance* [online], *16*(1), 293-308. Available at: https://doi.org/10.1111/rego.12318

Baudrillard, J., 1994. *Simulacra and simulation*. Trans.: S. F. Glaser. Ann Arbor: University of Michigan Press.

Beckers, A., and Teubner, G., 2023. Responsibility for algorithmic misconduct: Unity or fragmentation of liability regimes? *Yale Journal of Law & Technology* [online], *25*(Special Issue), 76. Available at: https://yjolt.org/node/1069

Bender, E. M., *et al*., 2021. On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* [online]. New York: ACM, 610-623. Available at: https://doi.org/10.1145/3442188.3445922

Berde, É., *et al*., 2025. GDP per capita and human capital investment in five countries after exhaustion of the first demographic dividend. *Regional Statistics* [online], *15*(5), 908-929. Available at: https://doi.org/10.15196/RS150504

Bovens, M., 2007. Analysing and assessing accountability: A conceptual framework. *European Law Journal* [online], *13*(4), 447-468. Available at: https://doi.org/10.1111/j.1468-0386.2007.00378.x

Burrell, J., 2016. How the machine "thinks": Understanding opacity in machine learning algorithms. *Big Data & Society* [online], *3*(1). Available at: https://doi.org/10.1177/2053951715622512

Buzan, B., Wæver, O., and de Wilde, J., 1998. *Security: A new framework for analysis*. Boulder: Lynne Rienner.

Citron, D. K., 2008. Technological due process. *Washington University Law Review*, *85*(6), 1249-1313.

Coglianese, C., and Lehr, D., 2017. Regulating by robot: Administrative decision making in the machine-learning era. *Georgetown Law Journal*, *105*(5), 1147-1223.

Commodity Futures Trading Commission (CFTC), 2024. *CFTC charges former CEO of carbon credit project developer with fraud involving voluntary carbon credits* [online]. 2 October. Press release 8994-24. Available at: https://www.cftc.gov/PressRoom/PressReleases/8994-24

Council on Environmental Quality (CEQ), 2025a. Removal of National Environmental Policy Act implementing regulations (interim final rule). *Federal Register*, 90, 10610.

Council on Environmental Quality (CEQ), 2025b. *Permitting technology action plan* [online]. 30 May. Washington, DC: Executive Office of the President. Available at: https://www.whitehouse.gov/articles/2025/05/trump-administration-launches-permitting-technology-action-plan

Crawford, K., ed., 2021. *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence* [online]. New Haven: Yale University Press. Available at: https://doi.org/10.12987/9780300252392

Department for Science, Innovation and Technology (DSIT), 2025. *Government-built Humphrey AI tool reviews responses to consultation for first time in bid to save millions* [online]. Press release, 14 May. Available at: https://www.gov.uk/government/news/government-built-humphrey-ai-tool-reviews-responses-to-consultation-for-first-time-in-bid-to-save-millions

Directive 2011/92/EU of the European Parliament and of the Council of 13 December 2011 on the assessment of the effects of certain public and private projects on the environment, as amended by Directive 2014/52/EU. *Official Journal of the European Union*, L 26, 1-21; L 124, 1-18.

Dryzek, J. S., 2009. Democratization as deliberative capacity building. *Comparative Political Studies* [online], *42*(11), 1379-1402. Available at: https://doi.org/10.1177/0010414009332129

Dryzek, J. S., and Niemeyer, S., 2008. Discursive representation. *American Political Science Review* [online], *102*(4), 481-493. Available at: https://doi.org/10.1017/S0003055408080325

Edelman, L. B., 1992. Legal ambiguity and symbolic structures: Organizational mediation of civil rights law. *American Journal of Sociology* [online], *97*(6), 1531-1576. Available at: https://doi.org/10.1086/229939

Galaz, V., Metzler, H., Daume, S., Olsson, A., Lindström, B., and Marklund, A., 2023. *AI could create a perfect storm of climate misinformation* [online]. arXiv:2306.12807. Available at: https://doi.org/10.48550/arXiv.2306.12807

George, A. L., and Bennett, A., 2005. *Case studies and theory development in the social sciences*. Cambridge, MA: MIT Press.

Glucker, A. N., *et al*., 2013. Public participation in environmental impact assessment: Why, who and how? *Environmental Impact Assessment Review* [online], 43, 104-111. Available at: https://doi.org/10.1016/j.eiar.2013.06.003

Goffman, E., 1959. *The presentation of self in everyday life*. Garden City, NY: Doubleday Anchor.

Haider, J., *et al*., 2024. GPT-fabricated scientific papers on Google Scholar: Key features, spread, and implications for preempting evidence manipulation. *Harvard Kennedy School Misinformation Review* [online], *5*(5). Available at: https://doi.org/10.37016/mr-2020-156

Hilgartner, S., 2000. *Science on stage: Expert advice as public drama*. Stanford University Press.

International Law Association (ILA), 2016. *Study group on due diligence in international law: Second report* [online]. D. French, Chair; T. Stephens, Rapporteur. Available at: https://www.ila-hq.org/en/study-groups

International Law Commission (ILC), 2001. Draft articles on prevention of transboundary harm from hazardous activities, with commentaries. *Yearbook of the International Law Commission 2001* [online], Vol. II, Part Two. Available at: https://legal.un.org/ilc/guide/9_7.shtml

International Law Commission (ILC), 2016. Draft articles on the protection of persons in the event of disasters, with commentaries. *Yearbook of the International Law Commission 2016*, Vol. II, Part Two.

Krawiec, K. D., 2003. Cosmetic compliance and the failure of negotiated governance. *Washington University Law Quarterly*, *81*(2), 487-544.

Kreps, S., and Kriner, D. L., 2024. The potential impact of emerging technologies on democratic representation: Evidence from a field experiment. *New Media & Society* [online], *26*(12), 6918-6937. Available at: https://doi.org/10.1177/14614448231160526

Krieger, H., Peters, A., and Kreuzer, L., eds., 2020. *Due diligence in the international legal order* [online]. Oxford University Press. Available at: https://doi.org/10.1093/oso/9780198869900.001.0001

Kroll, J. A., *et al*., 2017. Accountable algorithms. *University of Pennsylvania Law Review*, *165*(3), 633-705.

Mahajan, S., 2026. *Participatory provenance as representational auditing for AI-mediated public consultation* [online]. arXiv:2604.20711. Available at: https://doi.org/10.48550/arXiv.2604.20711

Mansbridge, J., *et al*., 2012. A systemic approach to deliberative democracy. *In*: J. Parkinson and J. Mansbridge, eds., *Deliberative systems: Deliberative democracy at the large scale* [online]. Cambridge University Press, 1-26. Available at: https://doi.org/10.1017/CBO9781139178914.002

Meyer, J. W., and Rowan, B., 1977. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology* [online], *83*(2), 340-363. Available at: https://doi.org/10.1086/226550

Ministry of Information Technology and Telecommunication (MoITT), 2025. *National Artificial Intelligence Policy 2025*. [online]. Islamabad: Government of Pakistan. Available at: https://www.moitt.gov.pk/SiteImage/Misc/files/National%20AI%20Policy.pdf

Nadeem, O., and Fischer, T. B., 2011. An evaluation framework for effective public participation in EIA in Pakistan. *Environmental Impact Assessment Review* [online], *31*(1), 36-47. Available at: https://doi.org/10.1016/j.eiar.2010.01.003

Nadeem, O., and Hameed, R., 2008. Evaluation of environmental impact assessment system in Pakistan. *Environmental Impact Assessment Review* [online], 28(8), 562-571. Available at: https://doi.org/10.1016/j.eiar.2008.02.003

National Environmental Policy Act of 1969, 42 U.S.C. §4332, as amended by the Fiscal Responsibility Act of 2023, Pub. L. No. 118-5, §321.

New York State Office of the Attorney General (OAG), 2021a. *Attorney General James issues report detailing millions of fake comments, revealing secret campaign to influence FCC's 2017 repeal of net neutrality rules* [online]. Press release, 6 May. Available at: https://ag.ny.gov/press-release/2021/attorney-general-james-issues-report-detailing-millions-fake-comments-revealing

New York State Office of the Attorney General (OAG), 2021b. *Fake comments: How U.S. companies & partisans hack democracy to undermine your voice*. New York: Office of the New York State Attorney General.

Nixon, R., 2011. *Slow violence and the environmentalism of the poor*. Cambridge, MA: Harvard University Press.

Ollino, A., 2022. *Due diligence obligations in international law* [online]. Cambridge University Press. Available at: https://doi.org/10.1017/9781009053082

Paine, J., 2026. Due diligence, obligations to cooperate and to regulate private actors: Insights from three climate change advisory opinions. *Environmental Policy and Law* [online], *56*(1-2), 18-29. Available at: https://doi.org/10.1177/18785395251394330

Pakistan Environmental Protection Act 1997 (PEPA) (Act No. XXXIV of 1997) [online]. Available at: https://elaw.org/wp-content/uploads/archive/Law-PEPA-1997.pdf

Pakistan Environmental Protection Agency (Pakistan EPA), 2000. *Pakistan Environmental Protection Agency (Review of IEE and EIA) regulations, 2000*. Islamabad: Government of Pakistan.

Pasquale, F., 2015. *The black box society: The secret algorithms that control money and information* [online]. Cambridge, MA: Harvard University Press. Available at: https://www.hup.harvard.edu/books/9780674970847

Power, M., 1997. *The audit society: Rituals of verification*. Oxford University Press.

Raji, I. D., *et al*., 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* [online]. New York: ACM, 33-44. Available at: https://doi.org/10.1145/3351095.3372873

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). *Official Journal of the European Union* [online], L, 2024/1689. Available at: http://data.europa.eu/eli/reg/2024/1689/oj

Rouvroy, A., and Berns, T., 2013. Algorithmic governmentality and prospects of emancipation: Disparateness as a precondition for individuation through relationships? *Réseaux* [online], *177*(1), 163-196. Available at: https://doi.org/10.3917/res.177.0163

Saren, T., 2025. Using AI in NEPA review: Legal challenges and judicial scrutiny. *Environmental Law Reporter* [online], *55*(6), 10596. Available at: https://www.elr.info/node/331794

Schmidt, V. A., 2013. Democracy and legitimacy in the European Union revisited: Input, output and "throughput". *Political Studies* [online], *61*(1), 2-22. Available at: https://doi.org/10.1111/j.1467-9248.2012.00962.x

Suchman, M. C., 1995. Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review* [online], *20*(3), 571-610. Available at: https://doi.org/10.5465/amr.1995.9508080331

Tian, X., and Shi, D., 2025. Facilitating or inhibiting: A study on the impact of artificial intelligence on corporate greenwashing. *Sustainability* [online], *17*(5), 2154. Available at: https://doi.org/10.3390/su17052154

United Nations General Assembly (UNGA), 2024. *Protection of persons in the event of disasters*. UN Doc A/RES/79/128, 4 December.

United Nations Office of Legal Affairs (UN OLA), 2026. *United Nations Diplomatic Conference of Plenipotentiaries on Protection of Persons in the Event of Disasters: Background* [online]. Available at: https://legal.un.org/diplomaticconferences/pped/bg.shtml

Wachter, S., Mittelstadt, B., and Russell, C., 2018. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, *31*(2), 841-887.

Wardle, C., and Derakhshan, H., 2017. *Information disorder: Toward an interdisciplinary framework for research and policy making* [online]. Strasbourg: Council of Europe. Available at: https://www.coe.int/en/web/freedom-expression/information-disorder

Weiss, M., 2019. Deepfake bot submissions to federal public comment websites cannot be distinguished from human submissions. *Technology Science* [online], 2019121801. Available at: https://techscience.org/a/2019121801/

Case law

International Court of Justice (ICJ), 2010. *Pulp Mills on the River Uruguay (Argentina v. Uruguay)*, Judgment, ICJ Reports 2010, p. 14.

International Court of Justice (ICJ), 2025. *Obligations of States in respect of climate change*, Advisory Opinion (23 July) [online]. Available at: https://www.icj-cij.org/case/187/advisory-opinions

International Tribunal for the Law of the Sea (ITLOS), 2011. *Responsibilities and obligations of States with respect to activities in the Area, Advisory Opinion (Seabed Disputes Chamber, Case No. 17), ITLOS Reports 2011*, p. 10.

International Tribunal for the Law of the Sea (ITLOS), 2024. *Request for an advisory opinion submitted by the Commission of Small Island States on Climate Change and International Law*, Advisory Opinion (21 May) (Case No. 31) [online]. Available at: https://www.itlos.org/en/main/cases/list-of-cases/request-for-an-advisory-opinion-submitted-by-the-commission-of-small-island-states-on-climate-change-and-international-law-request-for-advisory-opinion-submitted-to-the-tribunal/

Seven County Infrastructure Coalition v. Eagle County, Colorado, 2025. No. 23-975 (US Supreme Court, 29 May).