



Operational Translucency: A Third Mode of Algorithmic Visibility in Platform Governance

Dr. Deepika Rani^{1,*}

¹ Advocate and Researcher, Supreme Court of India, New Delhi, India

* Corresponding author: drdeepikaraniadvocate@gmail.com

ARTICLE INFO

Article history:

Received: 23-07-2026

Received in revised form:
03-08-2026

Accepted: 04-08-2026

Available online: 05-08-2026

Keywords:

AI Act

Algorithmic accountability

Algorithmic transparency

Digital Services Act

Platform governance

Researcher access

ABSTRACT

This paper develops the concept of operational translucency: a third mode of algorithmic visibility distinct from both transparency and opacity in platform governance. Drawing on the post-transparency literature, the opacity tradition, and emerging European regulatory primary sources, the paper argues that what platforms and regulators are now jointly building is neither full transparency nor total opacity, but a stable third regime characterised by partial, structured, asymmetric visibility in which different stakeholders see different things at different resolutions, on schedules and through interfaces controlled by the platform. Four dimensions are constitutive: resolution (granularity of disclosed information), audience (stakeholder-stratification), temporality (cadence and platform-controlled scheduling), and mediation (the interface and infrastructure of disclosure). The construct is positioned against qualified transparency, platform observability, and earlier passing uses of translucence. Implications are drawn for regulatory enforcement, audit practice, and researcher access, and a four-line research agenda is proposed. Naming the regime is the prerequisite for governing it.

1. INTRODUCTION

On 24 October 2025, the European Commission issued press release IP/25/2503, announcing that “the European Commission preliminarily found both TikTok and Meta in breach of their obligation to grant researchers adequate access to public data under the Digital Services Act” [1]. The proceedings followed earlier actions opened against X and AliExpress on the same Article 40 grounds. They sit within a regulatory architecture that simultaneously requires Very Large Online Platforms to publish transparency reports (DSA Articles 15 and 42), maintain a Europe-wide database of moderation actions (Article 24), permit independent audits (Article 37), run advertising repositories (Article 39), publish recommender-system parameter information (Article 27), and grant graduated data access to two distinct tiers of researchers (Article 40(4) and Article 40(12)).

What is striking about this regulatory complex is that no single stakeholder sees the platform whole. End-users see one slice; the public sees another; vetted researchers see a third; the wider research community sees a fourth; regulators and auditors see a fifth. Each slice has its own resolution. Each is delivered on its own schedule. Each pass through an interface (an API, a clean room, a transparency report, a Data Access Portal) that the platform, often in coordination with the regulator, controls.

This is not transparency in any meaningful sense Heald or Hood would recognise: no observer is positioned to scrutinise the system in its entirety [2-3]. Neither is it opacity in Pasquale’s sense: substantial information now flows to multiple audiences under legal compulsion [4]. It is something else, and that something else has no settled name. The literature treats it variously as a degraded form of transparency, as opacity in retreat, or as a temporary regulatory compromise. This paper argues that it is none of those. It is

a stable, distinct, and increasingly dominant *mode* of algorithmic visibility, and the absence of a theoretical vocabulary for it has consequences. It leaves regulators reasoning by analogy from older administrative-law traditions; it leaves critical scholars holding platforms to standards no one actually applies; and it incentivises platforms to perform transparency rather than enable accountability.

The construct this paper proposes is *operational translucency*: the platform-governance regime in which visibility into algorithmic systems is partial, structured, and asymmetric: different stakeholders see different things, at different resolutions, on schedules and through interfaces controlled by the platform (often in coordination with the regulator), with the result that no stakeholder sees the system whole, and the platform retains effective epistemic privilege over its own operations. Four dimensions are constitutive: resolution, audience, temporality, and mediation [5-7].

The paper makes three contributions. First, it names a governance regime that the existing vocabulary leaves nameless. Second, it specifies the regime's structure along four jointly necessary dimensions, distinguishing it from the closest neighbouring constructs in the literature: Pasquale's qualified transparency, Heald's asymmetric transparency, Hansen and Flyverbom's calibration of knowledge, Rieder and Hofmann's platform observability, and Birchall's passing use of "translucence". Third, it shows how the construct illuminates current regulatory primary sources (the DSA, the AI Act, and the audit and researcher-access ecosystem) in ways the transparency/opacity binary cannot. Section 2 surveys the transparency and opacity literatures and the adjacent constructs that almost name operational translucency. Section 3 defines the construct and specifies its four dimensions [8]. Section 4 reads three live regulatory regimes through it. Section 5 develops implications and a research agenda. Section 6 acknowledges limitations. Section 7 concludes.

2. THEORETICAL FOUNDATIONS

2.1 The transparency ideal and its critics

The transparency ideal in contemporary algorithmic governance descends from a long lineage in public administration, codified in Hood and Heald's edited *Transparency: The Key to Better Governance*. Heald's chapter "Varieties of Transparency" remains the canonical taxonomy: transparency operates in four *directions* (upwards, downwards, outwards, inwards); is either *event-based* or *process-based* in object; either *retrospective* or in *real time*; and either *nominal* or *effective*. Critically for the present argument, Heald already identifies *asymmetric* transparency (visibility flowing unidirectionally in a hierarchy, such that superiors observe subordinates without reciprocal visibility) as the hierarchical norm rather than the exception [9].

Heald's typology distinguishes transparency *directions* (upwards (subordinates to superiors), downwards (rulers to ruled), outwards (insiders to outsiders), and inwards (outsiders to insiders)), *objects* (event-based versus process-based), and *temporal modes* (retrospective versus real-time). These categories matter for the present argument because operational translucency selectively activates particular combinations along each axis simultaneously: a DSA Article 27 disclosure to end-users is downwards, event-based, on-demand; a DSA Article 37 independent audit is outwards-to-inwards, process-based, retrospective; an Article 39 ad repository is outwards, event-based, near-real-time. Heald's framework supplies the analytical primitives but does not yet constitute the construct: the regime emerges when a *single algorithmic system* is subjected to multiple Heald-type configurations simultaneously, with no observer in a position to integrate them.

For algorithmic systems, the binary "see/do not see" framing was sharpened by Pasquale's. The *Black Box Society*. Pasquale himself, however, was sceptical of total disclosure, proposing in his 2010 *Northwestern University Law Review* article "Beyond Innovation and Competition" a notion of *qualified transparency*: disclosure to designated third parties (auditors, regulators, vetted researchers) under confidentiality conditions [5]. This is the most important pre-existing construct neighbouring our proposal, and §2.3 returns to it.

Gillespie's characterisation of platforms as "custodians of the internet" and Klonick's account of platforms as "new governors" of online speech provide the institutional context for the present argument: the entities producing translucent surfaces are not merely commercial firms but quasi-public infrastructures whose disclosure practices constitute the empirical material of contemporary digital constitutionalism [10-12]. It is precisely because platforms have become governmental in function that the transparency they offer matters in the way that public-sector transparency once did, and it is precisely because they remain private firms with commercial and trade-secret interests that the transparency they offer can never be of the public-sector kind.

The decisive critical intervention against the transparency ideal in the algorithmic context is Ananny and Crawford’s “Seeing Without Knowing” in *New Media & Society*. Their argument enumerates ten limitations of transparency: it can be disclaimed, can produce harm, generates false binaries, privileges seeing over understanding, and treats algorithmic systems as bounded objects rather than as assemblages of “institutionally situated code, practices, and norms”. Their conclusion is that transparency, even if achieved, would not deliver accountability.

Fenster widens the diagnosis to disclosure regimes generally: the same belief that government information can be released on command also underwrites the belief that it can be kept secret on command, and both founder on the fact that information neither obeys disclosure rules cleanly nor stays put [15, 16]. The shared finding across this literature is that transparency, naively applied, becomes *transparency theatre*: a performance of disclosure that does not enable accountability. What this critical literature does not do, however, is name the regime that is actually emerging in transparency’s place.

2.2 The opacity literature

Parallel to the transparency critique runs the opacity literature, foundationally Burrell’s “How the Machine ‘Thinks’” in *Big Data & Society*. Burrell distinguishes three forms of algorithmic opacity: *intentional secrecy* (corporate or state withholding); *technical illiteracy* (the gap between code and reader); and *intrinsic opacity* of high-dimensional machine-learning models. Each form has its own remedy logic, and these remedies are not interchangeable [17].

Burrell’s typology pre-empted much of the explainable-AI debate that followed. Wachter, Mittelstadt and Floridi showed that the GDPR contains no enforceable right to explanation, only a “right to be informed” under Articles 13–15 and an ambiguous Article 22 [18]. Selbst and Barocas deepened the diagnosis by distinguishing *inscrutability* (we cannot read the rules) from *non-intuitiveness* (we can read them but they do not match human intuition) [19]. Rudin issued the strongest dissent from the explainable-AI consensus, arguing that high-stakes systems should be built from inherently interpretable models in the first place rather than retrofitted with post hoc explanations [20]. Edwards and Veale argued that the right-to-explanation discourse misallocates effort: individuals need systemic remedies (audits, ombuds, collective redress), not bespoke decision-by-decision explanations. Once again, the literature names what is insufficient about opacity, not what regime is replacing it [21].

2.3 Adjacent constructs and the novelty defence

The construct this paper develops is novel, but it sits in a thick neighbourhood that must be navigated explicitly. Table 1 summarises the differentiation.

Table 1. Differentiation of operational translucency from adjacent constructs

Construct	Origin	Locates visibility in	Distinguishing feature relative to operational translucency
Transparency (naive)	[9]	Total disclosure to all	Treats audience as unitary; operational translucency stratifies
Opacity	[4, 17]	Non-disclosure	Treats visibility as binary; operational translucency is graduated
Qualified transparency	[4, 5]	Privileged third parties	Dyadic disclosure; operational translucency is multi-audience, multi-resolution
Asymmetric transparency	[2]	Hierarchical flows	Direction-only; operational translucency adds resolution, timing, and mediation
Calibration of knowledge	[6]	Disclosure devices	Meta-analytic frame; operational translucency is the substantive regime
Platform observability	[7]	Outputs of socio-technical processes	Regulatory desideratum; operational translucency names the emergent regime
Strategic opacity	[22]	Purposeful obfuscation	A tactic within translucent regimes, not their description
Translucent databases	[23]	Cryptographic database design	Technical pattern at data layer, not a governance regime
Operational translucency	This paper	Stratified, scheduled, mediated visibility	Joint specification across four dimensions

Several entries warrant brief unpacking. Pasquale’s *qualified transparency* is the closest competitor and is therefore extended, not displaced, by the present construct. Operational translucency adds three features Pasquale does not theorise: it is multi-stakeholder rather than dyadic; it is multi-resolution rather than binary (open/closed); and it foregrounds *temporal* control by the platform.

Hansen and Flyverbom argue that “mediating technologies, conceptualized here as disclosure devices, have distinctive organizing properties ... They play a central role in attempts to shed light on objects, subjects and practices” . Their move is meta-analytic: naming what all transparency regimes do. Operational translucency is the substantive, descriptive name for one particular calibration architecture now dominant in platform governance.

Rieder and Hofmann propose *platform observability* as an alternative to transparency: regulators should focus on outputs of socio-technical processes rather than internal mechanisms. This is the closest sibling construct in the platform-governance literature [7]. The present construct differs in naming not a regulatory *desideratum* (what regulators should pursue) but the *actual emergent regime* (what regulators and platforms are jointly producing), which combines output observability with structured input and process disclosure under platform-controlled scheduling.

The word “translucence” itself has prior currency. Wayner’s *Translucent Databases* used it for a technical cryptographic pattern in which data is partly recoverable and partly hashed; this is a database-design idea, not a governance regime, and the borrowing here is metaphorical [23]. Birchall in her introduction to a *Theory, Culture & Society* special issue on secrecy and transparency, refers in passing to “the translucence of transparency” in extractive-industry disclosure, acknowledging an emerging in-between mode but not specifying it [8]. The present paper takes that observation as a starting point and develops it into a fully specified construct. Nissenbaum’s *contextual integrity* and Goffman’s presentation of self-frame the receiving end of operational translucency (different audiences receive information appropriate to the norms governing their context) and are constitutive of the audience dimension developed in §3[24, 25,26].

3. DEFINING OPERATIONAL TRANSLUCENCY

Drawing the strands together, the construct is defined as follows:

Operational translucency is the platform-governance regime in which visibility into algorithmic systems is partial, structured, and asymmetric: different stakeholders see different things, at different resolutions, on schedules and through interfaces controlled by the platform (often in coordination with the regulator), with the result that no stakeholder sees the system whole, and the platform retains effective epistemic privilege over its own operations.

Figure 1 positions the construct in the conceptual landscape, projecting two of the four dimensions (resolution and audience breadth) and plotting concrete DSA and AI Act mechanisms as illustrations of the multi-point, stratified visibility characteristic of operational translucency.




Figure 1. Operational translucency in the resolution × audience-breadth plane. The figure projects two of the four constitutive dimensions of operational translucency, situating the construct against the four adjacent regimes that occupy the corners of the conceptual space: naive transparency (empirically unrealised), opacity [4, 17], qualified transparency [4, 5], and performative transparency or “transparency theatre”. Concrete DSA and AI Act mechanisms are plotted to illustrate how operational translucency manifests as a multi-point footprint rather than a single position. Temporality (platform-controlled scheduling) and mediation (disclosure infrastructure) are the two further dimensions not depicted.

Four dimensions are constitutive.

3.1 Resolution

Resolution is the *granularity* of disclosed information. An end-user sees that “this post was removed for hate speech”. A low-resolution disclosure that names the *category* of moderation outcome but not its evidentiary basis. A vetted researcher under DSA Article 40(4) sees sampled moderation data with labels, timestamps, content categories, and selected metadata: a medium-resolution disclosure that supports descriptive and inferential analysis but does not expose the trained classifier or its training data. A regulator under Article 37 audit may see model cards, training-data lineage, threshold parameters, and red-team results: a high-resolution disclosure that supports structural diagnosis but stops short of full code and weights. No actor sees the maximum-resolution surface (raw training data, current model weights, and live-traffic logs simultaneously) because that surface remains internal to the platform. The point is not that any particular regulatory choice of resolution is wrong, but that the regime is constituted by the *stratification* of resolutions: different audiences sit at different points on the granularity gradient, and the design of that stratification is itself the governance question.

3.2 Audience

Audience captures the *stakeholder-stratification* of visibility. Translucent regimes presuppose differentiated audiences with distinct legal entitlements and informational competences. The DSA enumerates at least six: end-users (Articles 14, 17, 27); the public (Articles 24, 39, 42); vetted researchers (Article 40(4)); the wider research community (Article 40(12)); the Commission and Digital Services Coordinators (Articles 34, 37, 41); and, where relevant, business users. The AI Act parallels this with

deployers (Article 13), affected individuals (Articles 50, 86), and market-surveillance authorities. Each audience receives a different translucent surface designed for its informational role. The dimension owes its conceptual shape to two strands of prior work: Klonick’s account of platforms as “new governors” whose disclosure practices construct juridical audiences, and Nissenbaum’s contextual integrity, which establishes that appropriate information flow is audience-relative and that the same disclosure can be governance-positive in one context and norm-violating in another [25]. Goffman’s presentation of self and boyd’s work on networked audience-segregation supply the underlying sociological grammar [26-27].

3.3 Temporality

Temporality is the *cadence* and *platform-controlled scheduling* of disclosure. This is the most under-theorised dimension in the existing literature. Transparency reports are quarterly or annual; risk assessments are annual; advertising libraries lag real-time activity; researcher data-access permissions are granted in time-bounded windows; APIs are rate-limited and can be revoked. On 2 February 2023, Twitter (now X) announced that “starting February 9, we will no longer support free access to the Twitter API, both v2 and v1.1,” with Enterprise pricing reportedly starting at \$42,000 per month [28]. On 14 August 2024, Meta confirmed that “as of August 14, 2024, CrowdTangle is no longer available” [29]. Both actions unilaterally contracted real-time researcher visibility. Heald’s “real-time vs retrospective” distinction returns here in a sharper form: the *cadence* of disclosure is itself a governance variable, and one the platform typically sets.

3.4 Mediation

Mediation is the *interface and infrastructure* through which disclosure passes. Translucent disclosure does not occur in the open air; it passes through APIs, secure processing environments, data clean rooms, and mediated portals, including the DSA Data Access Portal mandated by the Article 40 Delegated Act adopted in July 2025. Hansen and Flyverbom’s point that disclosure devices have “distinctive organizing properties” applies here: the interface is not neutral [6]. It permits some queries and forbids others; it logs auditor access; it can be modified by the platform. The mediation dimension turns disclosure infrastructure into a governance object in its own right.

These four dimensions are *jointly* necessary. A regime with stakeholder-stratified disclosure but no resolution variance collapses into Pasquale’s qualified transparency. A regime with resolution variance but no scheduling collapses into Hansen and Flyverbom’s calibration of knowledge. A regime with both but no audience-stratification reduces to Rieder and Hofmann’s observability. Operational translucency is the conjunction of all four.

4. OPERATIONAL TRANSLUCENCY IN THREE REGULATORY REGIMES

4.1 The EU Digital Services Act

The Digital Services Act (Regulation (EU) 2022/2065) is the paradigmatic operational-translucency regime. Article 15 imposes baseline transparency-reporting obligations on all intermediary services; Article 24 strengthens these for online platforms; Article 27 requires user-facing recommender-system explanation in terms of “main parameters”; Article 37 mandates independent audits with reports submitted to the Commission; Article 39 requires advertising repositories; Articles 34 and 35 require risk assessments; Article 40 grants graduated data access, with Article 40(4) covering vetted researchers and Article 40(12) covering the wider research community; Article 42 imposes additional reporting on Very Large Online Platforms and Very Large Online Search Engines. Table 2 maps these surfaces onto the four dimensions of operational translucency.

Table 2. Translucent surfaces under selected DSA articles

Article	Audience	Resolution	Cadence	Mediation
Art. 15	Public	Low (aggregate)	Annual	Published reports
Art. 24, 42	Public	Low-medium (statements of reasons + VLOP metrics)	Quarterly / annual	DSA Transparency Database
Art. 27	End-users	Low (recommender “main parameters”)	On-demand	Interface controls
Art. 37	Commission, auditors	High (full audit access)	Annual	Audit reports

Article	Audience	Resolution	Cadence	Mediation
Art. 39	Public	Medium (ad metadata)	Near-real-time	Ad libraries
Art. 40(4)	Vetted researchers	Medium (non-public data)	Time-bounded windows	Data Access Portal
Art. 40(12)	Wider research community	Low-medium (public data)	Time-bounded	API or scraping permission

The DSA Transparency Database, launched in September 2023, illustrates the architecture in action. Kaushal analysed a representative sample of approximately 131 million statements of reasons submitted in a single month, finding substantial residual platform discretion over categorisation. Translucency is real; granularity is variable; categorisation logic remains with the platform [30].

Enforcement contests over the dimensions are now constant. The European Commission opened proceedings against X in July 2024 and accepted binding commitments from AliExpress in June 2025. On 24 October 2025, the Commission’s preliminary findings against TikTok and Meta contested two dimensions of operational translucency simultaneously for both platforms. The Commission preliminarily concluded that both TikTok and Meta had “introduced burdensome procedures and tools” for researchers to request access to public data: a *mediation* critique going to the design and usability of the disclosure interface, and that those tools resulted “frequently in partial or unreliable data”: a *resolution* critique going to the granularity and completeness of what is disclosed. The findings treat both dimensions as enforceable in their own right under Article 40, rather than as derivatives of a single underlying “transparency” obligation. The DSA Article 40 Delegated Act, adopted by the Commission in July 2025 and which entered into force on 29 October 2025 (five days after the preliminary findings) establishes the centralised Data Access Portal through which much of this access will be channelled, formalising the mediation dimension as a regulatory object.

These are not transparency disputes in the abstract; they are skirmishes over specific dimensions of an established translucent regime. The construct developed here makes the structure of the disputes legible.

4.2 The EU AI Act

The AI Act (Regulation (EU) 2024/1689) layers translucency by risk class. Article 13 mandates transparency to deployers of high-risk systems: instructions for use, capabilities, limitations, accuracy metrics, and foreseeable risks. Article 50 mandates transparency to end-users: chatbot disclosure, deepfake labelling, and biometric-categorisation notice. Article 86 establishes a right to explanation of individual decision-making for affected persons in certain high-risk cases. Articles 11, 12, and 18 require technical documentation and record-keeping accessible to market-surveillance authorities. The same high-risk hiring tool therefore generates four distinct translucent surfaces (one for the deployer, one for the end-user, one for the affected individual, and one for the regulator) each at its own resolution and on its own cadence. The Act codifies operational translucency rather than transparency; the only stakeholder positioned to see the system whole is the developer.

4.3 The audit and researcher-access ecosystem

Algorithmic auditing, descending from Sandvig, Hamilton, Karahalios and Langbort’s influential pre-conference paper, has now bifurcated along the audience dimension [31]. External audits (exemplified by the Algorithmic Justice League and by Mozilla’s RegretsReporter project on YouTube) operate from outside, often in legally precarious zones with low resolution and broad audience. Internal and contracted audits operate from inside with privileged access but constrained public reporting [32]. The DSA Article 37 regime formalises a contracted-external model. Costanza-Chock, Raji and Buolamwini [33] document the field’s central pathology: “claims that an AI product or system has been audited, whether by first-, second-, or third-party auditors, are difficult to verify and may potentially exacerbate, rather than mitigate, bias and harm”. This is a *translucency* pathology, not a transparency one: audits exist, reports are produced, but the verification chain is broken at the mediation layer.

The Social Science One partnership, launched on 11 July 2018 and co-chaired by Gary King and Nate Persily, illustrates a more acute pathology. The partnership was widely reported, including in published statements by its European Advisory Committee in December 2019, as having failed to deliver promised access: scholars received a “URL light data set” of substantially diminished scientific value relative to the original commitment, and the partnership’s difficulties became a paradigm case in subsequent platform-researcher access debates [34]. Translucent regimes can produce *negatively translucent* surfaces: visibility that is technically present but epistemically void.

The 14 August 2024 shutdown of CrowdTangle, despite the simultaneous launch of Meta’s Content Library, raised parallel concerns: the resolution and usability of the replacement tool were materially lower than those of the deprecated one, even as the formal claim of researcher access was maintained. The pattern across these cases (Social Science One, CrowdTangle, the X API contraction) is the same: when a translucent surface is contracted unilaterally by the platform, the legal claim to “transparency” can remain technically true while the substantive accountability function is hollowed out. The Coalition for Independent Technology Research and similar organisations have formed in part to contest these unilateral contractions, often by mobilising the resolution and mediation dimensions explicitly in their advocacy. The DSA Article 40 enforcement architecture is the regulatory response to exactly this dynamic: it is no longer enough for platforms to *claim* researcher access; the *resolution*, *mediation*, and *cadence* of that access have become independently enforceable.

5. IMPLICATIONS AND RESEARCH AGENDA

Naming operational translucency has four immediate analytic payoffs.

First, it permits a clearer normative debate. Critiques of “transparency theatre” presuppose that real transparency is achievable; critiques of “platform opacity” presuppose that opening the black box is the alternative. If translucency is the actual regime, the normative question shifts: which *resolutions*, *audiences*, *cadences*, and *interfaces* should the platform regulator pair set, and according to which principles? The argument is no longer about more or less disclosure but about the constitutional design of stratified disclosure. The Santa Clara Principles and Suzor’s work on digital constitutionalism supply candidate normative starting points, but neither yet engages the four-dimensional structure proposed here [35].

Second, the construct makes platform-strategic behaviour legible. Platforms compete to set resolution low, audience narrow, cadence retrospective, and interface theirs. Regulators push the other way. The DSA Article 40 enforcement actions of October 2025 become readable as contests over specific dimensions, not as abstract transparency failures. Industry submissions to consultations on the DSA Article 40 Delegated Act, on AI Act Article 50 implementation guidelines, and on the UK Online Safety Act’s transparency-reporting framework can be analysed as moves on the four dimensions: arguments for narrower vetted-researcher definitions are audience moves; arguments against detailed metadata in advertising repositories are resolution moves; arguments for longer reporting cycles are temporality moves; arguments for proprietary access infrastructure are mediation moves.

Third, the construct links platform governance to information-asymmetry economics and to disclosure-regulation theory in finance and consumer protection [36]. The distinctive feature is that the disclosing entity controls the disclosure infrastructure itself: a configuration without a clean analogue in earlier disclosure regimes. Securities disclosure runs through standardised forms filed with independent regulators; consumer-protection disclosure runs through legally specified label and notice requirements. Translucent platform disclosure, by contrast, runs through interfaces built and maintained by the entity being disclosed about, which is a structural feature meriting comparative-institutional analysis.

Fourth, the construct identifies an empirical research agenda. Four lines of work follow. *Comparative case studies* of how the four dimensions are negotiated across platforms and regulatory regimes: for example, comparing the DSA, the UK Online Safety Act 2023, and China’s algorithmic-recommendation regulation, the last of which exhibits a *state-led* translucency variant in which the audience–dimension reduces to the state plus a tightly constrained set of authorised observers. *Measurement studies* operationalising each dimension: resolution as a function of metadata fields disclosed; audience as the number and breadth of legally entitled observers; temporality as cadence and revocability; mediation as the cognitive and technical cost of using disclosure infrastructure. *Enforcement studies* tracing how regulators have learned to contest particular dimensions and what evidentiary architecture is required to enforce on each: work for which the October 2025 TikTok and Meta proceedings provide an unusually clean case. *Normative theorising* of how translucent regimes should be designed when full transparency is unavailable and opacity is unacceptable, drawing on Suzor’s digital constitutionalism, the Santa Clara Principles [35], and Nissenbaum’s contextual integrity for principles by which legitimate translucency can be distinguished from accountability theatre [25].

6. BOUNDARY CONDITIONS AND LIMITATIONS

Three boundary conditions delimit the construct. First, operational translucency presupposes that *some* visibility is being produced under legal or social compulsion; in regimes of pure trade-secrecy with no

disclosure obligations, opacity remains the relevant frame. Second, the construct presupposes that *multiple stakeholders* have plausible claims to visibility; in private bilateral contracts where only the contracting parties matter, qualified transparency in Pasquale’s sense suffices. Third, the construct’s force is greatest in *consequential* algorithmic systems (content moderation at scale, automated decisions in employment and credit, generative-AI deployments affecting public discourse) where accountability stakes justify the analytical apparatus [5].

Three limitations of the present paper should be acknowledged. The construct is at this stage conceptual; the measurement programme sketched in §5 is a proposal, not validated. The choice of the word “translucency” is contested; alternatives include “structured visibility”, “graduated disclosure”, and “stratified transparency”. The argument for keeping “translucency” rests on the optical metaphor’s intuitive grasp of partial but-real visibility, on Birchall’s prior governance use of the term, and on the contrast with the dominant transparent/opaque binary; the choice should be tested by the field [8]. Finally, the regulatory analysis is European-centred; comparative treatment of the United States’ Algorithmic Accountability Act efforts (not yet enacted as of 2026), the UK Online Safety Act 2023, and the Chinese state-led translucency variants embedded in the 2022 algorithmic-recommendation and deep-synthesis regulations is left for future work.

7. CONCLUSION

This paper has developed *operational translucency* as a third-mode construct for platform governance. The argument is that the transparency-versus-opacity binary has outlived its analytic usefulness: what platforms and regulators are jointly producing, in the DSA, in the AI Act, and in the audit and researcher-access ecosystem, is neither full transparency nor total opacity, but a stable third regime characterised by partial, structured, asymmetric visibility along four constitutive dimensions: resolution, audience, temporality, and mediation. The construct extends Pasquale’s qualified transparency by adding stakeholder-stratification, resolution variance, and platform-controlled scheduling; it differs from Rieder and Hofmann’s observability by naming the actual emergent regime rather than a regulatory desideratum; it builds on Birchall’s passing observation of “the translucence of transparency” by specifying the regime structurally.

Naming the regime is the prerequisite for governing it. So long as scholars and regulators reason as if full transparency were the goal and opacity the failure mode, they will miss the contests that actually determine algorithmic accountability contests over which audience sees what, at what resolution, on whose schedule, through whose interface. The European Commission’s October 2025 enforcement actions against TikTok and Meta are not transparency disputes in the abstract; they are skirmishes over specific dimensions of an established translucent regime. The construct developed here is offered as the vocabulary in which those skirmishes, and the constitutional design choices they imply, can be analysed. Governing algorithmic systems that shape behaviour, influence decisions and mediate the future of professional and democratic work therefore raises precisely the questions of accountability, fairness and institutional design with which this analysis began [37]; related work on artificial-intelligence deployment in the technology sector documents comparable accountability gaps in organisational settings [38].

REFERENCES

- [1] European Commission. (2025, October 24). *Commission preliminarily finds TikTok and Meta in breach of their transparency obligations under the Digital Services Act* (Press release IP/25/2503). https://ec.europa.eu/commission/presscorner/detail/en/ip_25_2503
- [2] Heald, D. (2006). Varieties of transparency. In C. Hood & D. Heald (Eds.), *Transparency: The key to better governance?* (pp. 25–43). Oxford University Press. <https://doi.org/10.5871/bacad/9780197263839.003.0002>
- [3] Hood, C. (2006). Transparency in historical perspective. In C. Hood & D. Heald (Eds.), *Transparency: The key to better governance?* (pp. 3–23). Oxford University Press. <https://doi.org/10.5871/bacad/9780197263839.003.0001>
- [4] Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press. <https://doi.org/10.4159/harvard.9780674736061>
- [5] Pasquale, F. (2010). Beyond innovation and competition: The need for qualified transparency in internet intermediaries. *Northwestern University Law Review*, 104(1), 105–173. <https://doi.org/10.2139/ssrn.1686043>
- [6] Hansen, H. K., & Flyverbom, M. (2015). The politics of transparency and the calibration of knowledge in the digital age. *Organization*, 22(6), 872–889. <https://doi.org/10.1177/1350508414522315>
- [7] Rieder, B., & Hofmann, J. (2020). Towards platform observability. *Internet Policy Review*, 9(4), 1–28. <https://doi.org/10.14763/2020.4.1535>

- [8] Birchall, C. (2011). Introduction to “Secrecy and transparency”: The politics of opacity and openness. *Theory, Culture & Society*, 28(7–8), 7–25. <https://doi.org/10.1177/0263276411427744>
- [9] Hood, C., & Heald, D. (Eds.). (2006). *Transparency: The key to better governance?* Oxford University Press. <https://doi.org/10.5871/bacad/9780197263839.001.0001>
- [10] Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press. <https://doi.org/10.12987/9780300235029>
- [11] Klonick, K. (2018). The new governors: The people, rules, and processes governing online speech. *Harvard Law Review*, 131(6), 1598–1670. <https://harvardlawreview.org/print/vol-131/the-new-governors-the-people-rules-and-processes-governing-online-speech/>
- [12] Suzor, N. P. (2019). *Lawless: The secret rules that govern our digital lives*. Cambridge University Press. <https://doi.org/10.1017/9781108666428>
- [13] Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
- [14] Ananny, M. (2016). Toward an ethics of algorithms: Convening, observation, probability, and timeliness. *Science, Technology, & Human Values*, 41(1), 93–117. <https://doi.org/10.1177/0162243915606523>
- [15] Fenster, M. (2006). The opacity of transparency. *Iowa Law Review*, 91(3), 885–949. <https://scholarship.law.ufl.edu/facultypub/46/>
- [16] Fenster, M. (2017). *The transparency fix: Secrets, leaks, and uncontrollable government information*. Stanford University Press. <https://doi.org/10.1515/9781503602670>
- [17] Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>
- [18] Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99. <https://doi.org/10.1093/idpl/ixp005>
- [19] Selbst, A. D., & Barocas, S. (2018). The intuitive appeal of explainable machines. *Fordham Law Review*, 87(3), 1085–1139. <https://doi.org/10.2139/ssrn.3126971>
- [20] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215. <https://doi.org/10.1038/s42256-019-0048-x>
- [21] Edwards, L., & Veale, M. (2017). Slave to the algorithm? Why a ‘right to an explanation’ is probably not the remedy you are looking for. *Duke Law & Technology Review*, 16, 18–84. <https://doi.org/10.31228/osf.io/97upg>
- [22] Xu, S., Xie, K., & Burtch, G. (2025). An empirical study of strategic opacity in crowdsourced evaluations. *MIS Quarterly*, 49(3), 1205–1220. <https://doi.org/10.25300/MISQ/2024/17441>
- [23] Wayner, P. (2002). *Translucent databases: Confusion, misdirection, randomness, sharing, authentication, and steganography to defend privacy*. Flyzone Press.
- [24] Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–157. <https://digitalcommons.law.uw.edu/wlr/vol79/iss1/6/>
- [25] Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press. <https://doi.org/10.1515/9780804772891>
- [26] Goffman, E. (1959). *The presentation of self in everyday life*. Anchor Books.
- [27] boyd, d. (2008). *Taken out of context: American teen sociality in networked publics* [Doctoral dissertation, University of California, Berkeley]. <http://www.danah.org/papers/TakenOutOfContext.pdf>
- [28] Calma, J. (2023, February 2). Twitter just closed the book on academic research. *The Verge*. <https://www.theverge.com/2023/2/2/23582615/twitter-removing-free-api-academic-research-misinformation>
- [29] Meta. (2024, August 14). *CrowdTangle*. Meta Transparency Center. <https://transparency.meta.com/researchtools/other-data-catalogue/crowdtangle/>
- [30] Kaushal, R., van de Kerkhof, J., Goanta, C., Spanakis, G., & Iamnitchi, A. (2024). Automated transparency: A legal and empirical analysis of the Digital Services Act transparency database. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)* (pp. 1121–1132). Association for Computing Machinery. <https://doi.org/10.1145/3630106.3658960>
- [31] Sandvig, C., Hamilton, K., Karahalios, K., & Langbort, C. (2014, May). *Auditing algorithms: Research methods for detecting discrimination on internet platforms*. Paper presented at the Data and Discrimination preconference of the 64th annual meeting of the International Communication Association, Seattle, WA, United States. <https://social.cs.uiuc.edu/papers/pdfs/ICA2014-Sandvig.pdf>

- [32] Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT '20)** (pp. 33–44). Association for Computing Machinery. <https://doi.org/10.1145/3351095.3372873>
- [33] Costanza-Chock, S., Raji, I. D., & Buolamwini, J. (2022). Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)* (pp. 1571–1583). Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533213>
- [34] King, G., & Persily, N. (2020). A new model for industry–academic partnerships. *PS: Political Science & Politics*, 53(4), 703–709. <https://doi.org/10.1017/S1049096519001021>
- [35] Santa Clara Principles. (2021). *The Santa Clara Principles on transparency and accountability in content moderation* (Version 2.0). <https://santaclaraprinciples.org/>
- [36] Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488–500. <https://doi.org/10.2307/1879431>
- [37] Pandey, J. K., & Kumar, R. (2025). Governing the algorithmic agent: Confronting overt and covert challenges to justice and the future of work. *International Journal of Law Management & Humanities*, 8(4), 1974–1984. <https://doi.org/10.1000/IJLMH.1110648>
- [38] Lalitha Kumari, P., & Pandey, J. K., et al. (2021). Impact of artificial intelligence on employees’ performance in information technology sector. *AIP Conference Proceedings*, 2587(1). <https://doi.org/10.1063/5.0150403>