

Algorithmic Media Trials: Social Media Censorship, Sub Judice Prejudice, and the Article 19–21 Balance in India

Kompella Madhuri*, Prof. Archana Gadekar

The Maharaja Sayajirao University of Baroda

*Corresponding Author

DOI: <https://dx.doi.org/10.47772/IJRISS.2025.909000508>

Received: 28 September 2025; Accepted: 04 October 2025; Published: 16 October 2025

ABSTRACT

This article theorises “algorithmic media trials” to describe how platform recommender systems, trending modules, and automated moderation transform reportage on ongoing criminal proceedings into sustained, personalised prejudicial publicity that risks sub judice prejudice while eluding traditional safeguards of open justice and contempt law thresholds. Anchoring the analysis in the Article 19–21 dialectic, it reconstructs the doctrinal tests under Article 19(2), necessity, proportionality, and least-restrictive means and reads them alongside fair trial and reputation as facets of Article 21 to propose a calibrated framework for managing prejudicial publicity in the platform era.

It synthesises the postponement jurisprudence (including the “real and substantial risk of serious prejudice” standard and neutralizing devices) with recent guidance that online removal orders must be necessary and proportionate, having considered less restrictive alternatives such as demotion, interstitials, and contextualisation.

The paper distinguishes private moderation from state-directed suppression and evaluates due-process deficits in emergent takedown pipelines by contrasting Section 69A’s reason-giving, hearing, and review architecture with extra-statutory workflows that incentivise over-removal and undermine the constitutional predicates that saved blocking powers from invalidation. Against this backdrop, it proposes a platform-compatible postponement template, a “sub judice mode” for intermediaries privileging demotion over deletion with notice-and-appeal, and a 69A-compliant state protocol that rules out bulk, reasonless demands. Case studies of algorithmically amplified controversies illustrate why virality and persistence alter prejudice calculus and why timing thresholds in contempt doctrine must account for digital reach and recirculation. The contribution doctrinally integrates fair trial safeguards with algorithmic mechanics and translates that integration into operational guidelines that preserve robust reporting while minimizing trial prejudice, without reverting to blanket prior restraint.

Keywords: Media Trials, Proportionality, Sub Judice Prejudice, Algorithmic Amplification, Section 69 India

INTRODUCTION

“Trial by media” has long challenged courts’ duty to ensure fair adjudication while upholding open justice and press freedom. In the platform era, however, the mechanics of prejudice have shifted from editors and broadcast slots to model-driven, continuous feeds. This paper theorises “algorithmic media trials” (AMT): the transformation of reportage on pending criminal proceedings into sustained, personalized, and high-salience exposure by engagement-optimised ranking systems, trending modules, and automated moderation pipelines. AMT intensifies risks of sub judice prejudice while evading the temporal and scope thresholds that structured traditional contempt jurisprudence.

The central claim is twofold. First, the constitutional equilibrium between Article 19(1)(a) and Article 21 must be reconstructed for a platform society, where the main risk is not one-off publication but platform-enabled

amplification, persistence, and personalization. Second, the doctrinal lever to regulate AMT without chilling press and public oversight is a rigorous proportionality analysis under Article 19(2) that obliges courts and state actors [1] to adopt least-restrictive means tailored to algorithmic mechanics especially demotion and contextualization before ordering removal or blanket prior restraint.

Methodologically, the article integrates doctrinal analysis (Articles 19 and 21, contempt and postponement jurisprudence, necessary-and-proportionate takedown standards), regulatory architecture (Section 69A and the IT Rules), and platform mechanics (recommenders, virality, and content curation). It culminates in operational proposals: a platform-compatible postponement template, a court-recognised “sub judice mode” SOP for intermediaries, and a 69A-compliant protocol for state requests.

Theorising Algorithmic Media Trials

Defining AMT

Algorithmic media trials arise when platform optimisation models ranking functions, trending surfaces, notifications, and auto-suggest convert episodic coverage of ongoing cases into an always-on stream of prejudicial narratives. The platform is not a passive conduit; its architecture actively determines which items receive reach, in what contexts, and how repeatedly users are exposed. Because these black-boxed models are tuned for engagement, they are more likely to amplify sensational, affect-laden, or polarising content features that often characterise speculative commentary on ongoing criminal cases. Doctrinally, this necessitates a shift in locus: from the intent or negligence of the publisher to the foreseeable, structural effects of algorithmic distribution. The legal system must recognise that a platform’s architecture can create prejudicial publicity predictably, even if the original story is factually accurate or neutrally framed. A focus on the mechanism of amplification (not just the content) is thus essential.

The Persistence and Virality Calculus

Traditional contempt doctrine calibrated restrictions by reference to time-sensitive trial stages and one-off publications. In platform ecosystems, content persists, is recontextualised by search and recommendation, and is re-amplified during procedural milestones (arrest, charge-sheet, trial commencement). The risk is not fleeting, it is persistent and personalised. Filters and feedback loops ensure repeated exposure, increasing salience and anchoring impressions about guilt or character. This persistence erodes the utility of purely temporal safeguards and suggests judicial remedies must address amplification rather than only removal of specific URLs. [2]

The Article 19–21 Dialectic Reconstructed

Fair Trial and Reputation (Article 21)

Article 21 embraces fair procedure, an unbiased adjudicator, presumption of innocence, and the dignity-anchored right to reputation. Prejudicial publicity threatens these interests by shaping public and participant perceptions, potentially intimidating witnesses or placing subtle pressure on decision-makers. [3]

Any governance of AMT must therefore preserve the core of fair trial rights while recognising that due process applies to speech restrictions and to the procedures by which those restrictions are imposed.

Free Speech and Open Justice (Article 19(1)(a))

Open justice demands public oversight and the press’s ability to report on state power and judicial processes. The public’s right to know about judicial proceedings is a central strand of Article 19(1)(a). Yet, when open justice is mediated by algorithms that systematically amplify prejudicial narratives, doctrine must separate factual reporting and scrutiny from targeted, engagement-chasing amplification that undermines the adjudicative process. The constitutional task is not to roll back open justice but to calibrate its exercise in algorithmic environments.

Article 19(2) Proportionality and Necessity

Restrictions must be: (i) for a legitimate aim (e.g., integrity of justice administration), (ii) suitable, (iii) necessary (no equally effective, less-restrictive alternative), and (iv) proportionate *stricto sensu* (benefit outweighs rights cost). [4] This four-part test is especially apt in the digital context because platforms provide a continuum of interventions beyond binary deletion: labelling, interstitials, contextualisation, friction, downranking, and geofenced or time-bound measures. Courts must therefore meaningfully apply least-restrictive-means (LRM), not merely recite it. [5]

Proportionality for the Platform Age

The four-pronged standard

Firstly, with respect to aim and legitimacy, the objective of any restriction must be preventing a “real and substantial risk of serious prejudice” to a pending proceeding. This is not a vague or abstract concern but one tied directly to the constitutional guarantee of a fair trial under Article 21. The focus must be on preserving the integrity of fact-finding, ensuring witness autonomy, and safeguarding adjudicative independence. Importantly, the aim must be tethered to a specific proceeding and not to speculative harms, with courts identifying concrete pathways of prejudice such as the tainting of identification evidence, the chilling of witnesses, or the premature adjudication of guilt in the public domain.

Secondly, under suitability or rational connection, the measure adopted must demonstrate an evidence-based likelihood of reducing the identified risk within the specific platform context. It is insufficient for a court to merely assert that a restriction will work; it must instead articulate how and why a particular neutralising device such as the downranking of recommendations actually curtails the recirculation or amplification that produces prejudice. The suitability analysis thus demands a clear causal explanation between the chosen measure and the harm sought to be prevented.

Thirdly, in relation to necessity and the least-restrictive-means requirement, any intrusive measure such as deletion or blocking must only be considered after evaluating concrete digital alternatives that are calibrated to algorithmic mechanics. Courts and regulators should systematically consider less drastic tools contextualisation through interstitial warnings or official summaries, demotion mechanisms like exclusion from trending or downranking in search feeds, the introduction of friction through share prompts, and time-bound or geographically scoped restrictions. Only if these calibrated measures are demonstrated to be insufficient on the facts should stronger forms of suppression be invoked, and the reasons for rejecting lesser means must be explicitly recorded.

Finally, the test of balancing or proportionality *stricto sensu* requires weighing the fair-trial gains of the restriction against the aggregate costs imposed on free expression, public access to information, and the principle of open justice. This stage of the analysis must acknowledge that contemporaneous reporting and public scrutiny of judicial proceedings carry immense democratic value, especially in a system committed to transparency and accountability. The restriction can be justified only if the risk of trial prejudice within the particular procedural window is so substantial, and so well-documented, that it outweighs the speech and democratic costs of limiting public discourse. In other words, proportionality requires not abstract balancing but a reasoned and explicit evaluation of competing constitutional goods in the given context.

The LRM gap and judicial capacity

The least-restrictive-means (LRM) gap highlights a pressing challenge for judicial capacity. Too often, LRM is invoked as a conclusory label, without any demonstrable reasoning that alternatives were seriously tested. To avoid this, judicial orders must move from abstraction to proof: they should explicitly examine platform-specific levers such as trending exclusion, recommendation off-switches, search demotion, interstitial warnings, or rate limits and explain why these granular measures cannot adequately mitigate the identified risk within the necessary timeframe. This kind of record transforms LRM from rhetoric into a disciplined method.

Given the technical unfamiliarity of many benches, simple judicial aids are essential. These could take the form of bench-book checklists, standardised affidavits from neutral technical amici, and structured platform declarations [6]

With these tools, judges would be able to map the precise “prejudice vector” whether it flows through feeds, search results, trending algorithms, or push notifications then match it to an appropriate neutralising device. Crucially, courts should also set review dates to reassess whether the chosen intervention has actually proven effective, building an iterative process into the order itself.

The test must also be grounded in robust evidentiary inputs. Orders should not rest on speculation but draw upon concrete data: reach metrics, virality curves, prevalence of certain queries, engagement hot-spots, and patterns of amplification around sensitive procedural events like hearings or witness examinations [7]

In genuinely urgent phases, where real-time decisions are unavoidable, provisional measures may be justified but these must be accompanied by mandated post-hoc validation and rapid recalibration once better data becomes available. This ensures that necessity is not frozen in urgency but evolves with evidence.

Finally, a discipline of systematic record-keeping is indispensable. Judicial orders should capture not only the alternatives considered but also the specific reasons for rejecting each one, and the rationale for concluding that the chosen remedy is the narrowest effective option. Such documentation does more than safeguard appellate scrutiny; it creates an institutional memory, enabling future benches to learn from past calibration rather than beginning each case from scratch. Over time, this cumulative practice can transform LRM analysis from a site of weakness into a core strength of the judiciary’s response to digital-era trial prejudice.

Least-restrictive means and algorithmic hierarchy

A coherent approach to least-restrictive means in the algorithmic context requires a presumptive hierarchy of interventions that reflects both the gravity of prejudice risk and the corresponding cost to speech. At the base of this hierarchy lies contextualisation, which should ordinarily be the default first step. This includes interstitial warnings, neutral context boxes, official court notices or summaries, and “sub judice” labels attached to identified URLs, queries, hashtags, or handles. Contextualisation works because it preserves access and reportage while correcting misimpressions at the point of consumption, thereby dampening prejudicial inference without suppressing speech. Such measures should be deployed whenever there is a plausible risk of prejudice tied to a pending proceeding, and should be paired with time-bound review keyed to procedural milestones such as the examination of vulnerable witnesses.

Where contextualisation alone proves insufficient particularly in high-virality settings such as leaks of confessional material, pre-trial identification footage, or threats to witness anonymity the next step is calibrated demotion. This entails removing content from recommendations and trending surfaces, downranking it in feeds and search results, de-duplicating repetitive clips, or introducing friction in sharing. Demotion does not erase speech but targets the amplification mechanics that make content prejudicial, reducing its salience without extinguishing it. Its effectiveness can be measured through exposure and engagement deltas, providing courts with empirical grounds for necessity findings. At the apex lies deletion or blocking, which is an exceptional remedy. This should be reserved only for cases of direct and acute prejudice such as the publication of sealed evidence or the doxxing of protected witnesses and then only upon particularised findings that both contextualisation and demotion are inadequate. Because it eliminates speech and public access altogether, it must be tightly time-bound, precisely scoped, and subject to periodic judicial review.

To operationalise this hierarchy in practice, judicial orders must be crafted with specificity, reviewability, and transparency. This means enumerating exact identifiers, URLs, handles, hashtags, search queries rather than broad topics, and tying the measure to a concrete procedural phase and prejudice pathway. Orders should require platforms to report anonymised exposure metrics such as impressions, click-through rates, or trending status, and set short review hearings, typically within one to three weeks, to reassess whether the chosen measure suffices or must be adjusted. Tailoring should be the rule: measures should be time-bound, event-bound, and where appropriate, geo-bound, with automatic sunset clauses keyed to trial milestones. Transparency is equally

critical. While sensitive identifiers may be preserved under sealed annexures to prevent gaming, platforms should still be required to maintain de-identified transparency logs (noting case tag, measure type, and dates), to issue itemised notices to creators or affected parties, and to provide them with expedited appeal channels.

Finally, proportionality in this domain depends on a disciplined evidentiary foundation and institutional safeguards. Courts must insist on concrete evidence of risk and amplification content type, audience reach, virality curves, trending trajectories, cross-platform spillovers and ensure that any move to deletion is preceded by a record of contextualisation and demotion attempts, with reasons for their insufficiency. Orders must link amplification profiles to specific trial risks such as eyewitness memory contamination, juror prejudice (in jurisdictions with lay adjudicators), or intimidation of witnesses, avoiding generic invocations of “fairness.” At the same time, courts must avoid common pitfalls such as defaulting to deletion, adopting topic-level bans, or issuing timeless orders without review. Platforms, for their part, should be required to maintain pre-built “sub judice” playbooks for rapid deployment of contextualisation tools, measurement protocols for exposure and engagement, and dedicated escalation tracks for court-tagged cases with strict service-level targets. Together, these judicial and platform-side practices transform LRM from an abstract aspiration into a workable algorithmic discipline, balancing the imperatives of fair trial with the constitutional commitments to speech and open justice.

Postponement Jurisprudence and “Real and Substantial Risk”

The Sahara standard

The Constitution Bench in Sahara conceptualised postponement as a narrowly tailored “neutralising device,” not a speech penalty, available only upon a clear, particularised showing of a “real and substantial risk of serious prejudice” to the fairness of a pending proceeding and the further finding that “reasonable alternative methods will not prevent the risk.” The standard embeds a structured proportionality test: the applicant bears the burden to displace the default of open justice by proving an identifiable prejudice pathway (e.g., tainting eyewitness memory, contaminating identification evidence, intimidating witnesses, or pressure on adjudicators), the immediacy and gravity of that risk within the procedural posture, and the insufficiency of lesser measures tailored to that posture. The order must be temporally cabined, content-specific, and reviewable; it functions as preventive relief keyed to trial integrity, not content censorship or viewpoint discrimination. Critically, Sahara’s logic is remedial, not punitive: it demands a reasoned inquiry into alternatives and forbids broad, topic-level restraints, thereby securing Article 19(1)(a) and the public’s right to know even as it vindicates Article 21’s fair trial mandate through precise, time-bound relief. [8]

Updating neutralising devices for AMT

In platform environments, “reasonable alternatives” must be understood as digital, mechanism-targeted neutralisers that directly address algorithmic amplification, the true vector of prejudice risk in AMT. This requires courts to adopt a three-step ladder consonant with Sahara’s structure: first, establish the real and substantial risk with platform-specific evidence (reach/virality curves, trending status, search-query prevalence, timing around sensitive evidentiary phases); second, test least-restrictive digital measures that neutralise amplification rather than speech, contextualisation (interstitials, neutral court summaries, sub judice labels) and demotion (exclusion from recommendations/trending, feed and search downranking, de-duplication, share friction); third, reserve deletion or full postponement for the residual class of items where granular findings show that contextualisation and demotion cannot adequately avert serious prejudice (for example, publication of suppressed confessions, protected witness identities, sealed exhibits). Orders should specify identifiers (URLs, handles, hashtags, search queries), set calibrated durations tied to procedural milestones, mandate platform transparency logs and metrics for review (exposure and engagement deltas), and include expedited appeal channels for affected speakers. By recentring “reasonable alternatives” on digital LRM, courts preserve contemporaneous reporting and open justice while mitigating the precise amplification mechanics that produce sub judice prejudice, thus carrying Sahara’s doctrinal balance forward into the algorithmic age.

Due Process and Digital Censorship: Section 69a vs Extra-Statutory Workflows

Section 69A of the IT Act and the 2009 Blocking Rules survived constitutional scrutiny [9] because they embedded procedural guardrails, reason-giving, opportunities for hearing, and review by a competent authority.

These predicates of specificity, accountability, and ex post supervision were sufficient to satisfy Articles 14, 19, and 21. In other words, Section 69A's legitimacy derives not from the breadth of the power it confers but from the due process scaffolding that ensures it is exercised in a reasoned, accountable, and reviewable manner.

Yet in practice, emergent workflows increasingly bypass this constitutionally significant framework. Bulk portal requests, templatised notices without reasons, or informal directions from low-rank officials place platforms in a position of compliance by fear rather than law. When takedowns or suppressions occur under threat of liability rather than pursuant to the structured mechanism of Section 69A, private moderation becomes indistinguishable from state-pressured censorship. This circumvention erodes the very safeguards that allowed Section 69A to survive constitutional challenge and turns digital regulation into a field of opaque executive action.

Compounding the problem are the safe-harbour pressures built into the IT Rules. Platforms, fearing loss of immunity, are incentivised toward binary deletion as the safest route, even when less speech-restrictive measures would suffice [10]

This regulatory environment sits uneasily with judicial preference for least-restrictive means, particularly contextualisation and demotion. If constitutional values are to be reconciled with platform incentives, the law must explicitly recognise that these calibrated responses constitute compliance for safe-harbour purposes. Otherwise, the perverse effect is to encourage over-removal, undermining both proportionality and the public's right to information.

Operational Guidance

To realign practice with constitutional principle, courts and platforms alike need structured operational guidance. A "sub judice mode" standard operating procedure would allow platforms, on a court's direction identifying a sensitive proceeding and content vector, to apply calibrated measures. Contextualisation can be achieved by interstitials and court-vetted summaries attached to flagged URLs, hashtags, and queries, paired with a "sub judice" advisory. Demotion would then remove such items from recommendation systems, trending modules, or auto-play chains, while downranking them in feeds and search. These interventions should be accompanied by notice and appeal channels for creators, and transparency logs recording the case identifier, measure type, and duration.

Courts, for their part, can adopt a postponement template to ensure proportionality and precision. Orders should clearly articulate the legal basis and the specific prejudice risk identified, list precise identifiers rather than broad topics, and require tiered remedies, contextualisation and demotion first, blocking only if lesser measures fail. Time-boundedness should be mandatory, linking the restraint to procedural stages such as the completion of a witness examination, with review hearings and automatic sunset clauses unless renewed on fresh findings. Such templates not only discipline judicial reasoning but also provide predictability for platforms implementing them.

The executive too must adhere to 69A-compliant request protocols. Every request should be particularised, justified in writing with reference to grounds under Article 19(2), and, wherever feasible, allow for intermediary or user submissions before review by the competent authority. Bulk or blanket demands must be avoided; if multiple identifiers are implicated, each must carry its own reasons. Crucially, the framework must align liability by recognising that court-directed demotion or contextualisation constitutes full compliance, preventing coercion of platforms into over-removal. Only by insisting on predicate fidelity, particularisation, and liability alignment can executive action stay within constitutional bounds while safeguarding both fair trials and open justice.

CONCLUSION

Algorithmic media trials reveal a doctrinal lag between traditional contempt jurisprudence and the realities of platform-mediated publicity. Prejudicial influence today is not the product of isolated broadcasts but the foreseeable outcome of engagement-optimised distribution systems that amplify, persist, and personalise coverage of pending proceedings. Doctrinal fidelity to the Article 19–21 balance requires that proportionality and least-restrictive means [11] be applied not to abstract speech categories but to the mechanics of algorithmic amplification.

The calibrated framework proposed here demonstrates that protecting fair trials need not entail retreat from open justice. Courts can demand neutralising devices contextualisation and demotion before considering exceptional remedies of deletion or postponement. The Sahara standard, reinterpreted for the platform era, thus preserves contemporaneous reporting while ensuring that virality and persistence do not corrode adjudicative integrity. For this equilibrium to hold, however, due process predicates must be preserved. Executive action must remain tethered to the reason-giving and review architecture that legitimised Section 69A, resisting the drift toward opaque portal-driven suppression. Platforms, in turn, require regulatory incentives that recognise proportionate LRM, contextualisation and demotion as full compliance for safe harbour purposes, rather than coercing them into blunt deletion.

Ultimately, the sub judice mode SOP, the postponement template, and the 69A-compliant protocol translate doctrine into practice. They operationalise proportionality in a way that allows robust reporting, public oversight, and digital transparency to coexist with trial integrity. The task is not to resurrect blanket prior restraint but to craft remedies that are narrow, reviewable, and digitally literate ensuring that the constitutional dialectic between free speech and fair trial endures in the algorithmic age.

REFERENCES

1. Maneka Gandhi v. Union of India, 1 SCC 248 (1978).
2. Gillespie, T. (2018). *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press.
3. O'Hara, E. R. (2020). Extrajudicial statements and prejudice in the digital age. *William & Mary Law Review*, 61(3), 1105–1168.
4. Barak, A. (2012). *Proportionality: Constitutional Rights and their Limitations*. Cambridge University Press.
5. Sahgal, R. (2024). *Proportionality Review and Economic and Social Rights in India*. University of Birmingham (working paper).
6. Custers, B. (2024). A fair trial in complex technology cases: Why courts and judges need technical literacy. *International Review of Law, Computers & Technology*, 38(1–2), 1–25.
7. Bhupatiraju, S., Chen, D. L., & Kapur, D. (2020). The promise of machine learning for the courts of India. *National Law School of India Review*.
8. Kharak Singh v. State of U.P., AIR 1963 SC 1295.
9. Shreya Singhal v. Union of India, 5 SCC 1 (2015).
10. Jain, S. (2015). *Sahara India Real Estate Corp. Ltd. v. SEBI: Balancing regulation and constitutional constraints* [SSRN].
11. Justice K.S. Puttaswamy (Retd.) v. Union of India, 10 SCC 1 (2017).