

AI-Generated Consumer Profiles, Transparency and Fairness: Consumer Protection in India's Digital Marketplace

Aman Gautam*¹, Bhupinder Singh²

¹Research Scholar, Sharda School of Law, Sharda University, Greater Noida,

Email Id: amn.gtm@gmail.com

²Professor, Sharda School of Law, Sharda University, Greater Noida,

Email Id: Bhupinder.singh@sharda.ac.in

***Corresponding Author**

Aman Gautam

Research Scholar, Sharda School of Law, Sharda University, Greater Noida,

Email Id: amn.gtm@gmail.com

ABSTRACT

Consumers increasingly encounter decisions shaped by profiles created from their personal information. A person may provide information to one business without knowing that an analytics supplier will use it to produce a score or prediction that affects access to a loan, job-related service or other commercial opportunity. The problem becomes harder to address when the business making the decision relies on a third-party system and the supplier treats its methods as confidential. This conceptual legal article examines how Indian data protection law can support meaningful consumer understanding and challenge without requiring the disclosure of an entire proprietary model. It uses doctrinal analysis of the Digital Personal Data Protection Act, 2023 and the notified Rules of 2025, alongside selected European legislation, judicial decisions and scholarship. The article connects three consumer-facing concerns: transparency about the profile used, a fair opportunity to question its relevance or accuracy, and correction that reaches the businesses still relying on it. It proposes a proportionate framework based on a decision record, a specific justification for secrecy, independent review where legally authorised, and documented follow-through after a successful challenge. The article does not report survey, interview or audit findings, and its proposals are normative rather than empirical. Its central argument is that consumer protection should focus not only on access to personal data, but also on whether a person can understand and meaningfully challenge the commercial use of an AI-generated profile.

Keywords: AI-generated profiles; consumer privacy; transparency; algorithmic fairness; data protection

1. INTRODUCTION

Consider a consumer who is refused a financial service after a company receives an automated assessment from an external analytics provider. The consumer asks what information led to the refusal. The company supplies the details originally entered, but does not explain the assessment that influenced the decision. The supplier says that its scoring method is confidential, while the company says that it does not control the underlying model. Neither response establishes whether the profile was accurate, relevant or fairly used. This is a hypothetical example, not a report of an observed dispute. It illustrates a consumer-protection problem that can arise when the business facing a person and the business able to explain a profile are different. The key question is how a consumer can understand and challenge a consequential profile when responsibility is divided between businesses.

The Digital Personal Data Protection Act, 2023 supplies the principal statutory framework examined here.¹ Its commencement must be distinguished from its enactment. The notification dated 13 November 2025 stages the operation of different provisions, with the principal processing duties and the rights in sections 11 to 17 assigned to the phase eighteen months after Gazette

publication.² The notified Digital Personal Data Protection Rules, 2025 likewise contain a staged commencement arrangement.³ The Rules are read together with the corrigenda dated 10 December 2025.⁴ This article consequently analyses the enacted architecture and its scheduled implementation, rather than assuming that every substantive entitlement is already operational on the research date of 10 October 2026. That distinction matters because a proposal for implementation guidance is different from a description of an entitlement that an individual can presently enforce. It also creates a useful opportunity to consider how organisations should document inferences before systems and contracts make later review unnecessarily difficult.

The central argument is that formal privacy rights may offer limited practical value if a consumer cannot identify the profile used, understand its role in a decision or obtain an effective correction. A personal inference and the model used to generate it are different objects of legal analysis. Disclosing the inference need not require disclosure of the entire model; equally, disclosing the inference alone may not reveal a wrong input, an unsuitable comparison or continued use of an outdated assessment. The article therefore proposes a consumer-

focused process in which the business using the profile remains answerable for its commercial decision, relevant suppliers preserve evidence needed for review, and a successful correction is followed through to relevant recipients. This approach connects transparency with practical redress instead of treating explanation as an end in itself.

2. RESEARCH QUESTION AND METHOD

The research question is how Indian law should allocate transparency and correction responsibilities when an AI-generated profile about a consumer is produced and used across a commercial supply chain. Three subsidiary questions guide the analysis. First, how should the law distinguish the personal profile used in a consumer-facing decision from confidential information about the system that produced it? Second, what do the Indian Act and Rules currently provide for access, accuracy and correction? Third, what legal and contractual arrangements could help ensure that a consumer's successful challenge reaches both the supplier and the business that made the decision? For this article, a consequential consumer profile means an assessment linked to an identifiable person that materially influences access to a service, an opportunity or another significant commercial outcome. This is an analytical term proposed for the article, not a defined category in the Act.

The method is doctrinal and comparative. Primary materials include the Indian Act, its commencement notification and notified Rules, the General Data Protection Regulation, the European Trade Secrets Directive, the TRIPS Agreement and the judgment in *Dun and Bradstreet Austria*. Selected scholarship is used to locate the argument within existing debates about inference, explanation and confidentiality. Sources were located through targeted searches for the named instruments, decisions and concepts, with preference given to official legal repositories, journal publications and institutional research archives. This is a focused examination rather than a systematic literature review. It makes no claim to have retrieved every publication on the subject. The legal analysis separates the wording of enacted provisions, arguments about their interpretation and recommendations for amendment. Hypothetical applications test the internal coherence of the proposal, but they cannot establish its effectiveness in practice or the frequency of the harms discussed.

Consumer-research relevance. The article approaches transparency as a condition for consumer agency rather than as a demand for technical disclosure for its own sake. A consumer's ability to make sense of a profile may affect whether the person can question a decision, correct an error or decide whether to continue dealing with a business. Fairness is treated here as a normative design concern: consumers should have a meaningful route to challenge materially consequential uses of their information. The article does not measure trust, perceived fairness or purchasing behaviour, so it does not claim that the proposed safeguards have been shown to improve those outcomes. Instead, it identifies legal conditions that future empirical consumer research could test.

The unit of analysis is the path followed by one profile from its creation to its use in a consumer-facing decision

and, where challenged, to its correction. This makes a practical problem visible: a consumer may correct information held by one business while an old score remains active elsewhere. The article evaluates its proposal against four normative criteria: whether the consumer receives useful information, whether a competent reviewer can examine supporting evidence, whether legitimate confidentiality is protected, and whether a successful challenge changes the records and uses that caused the problem. These criteria are analytical standards, not findings from interviews, surveys or an audit of commercial systems.

3. LITERATURE AND CONCEPTUAL CONTRIBUTION

The problem of regulating inferences is established in privacy scholarship. Wachter and Mittelstadt argue for a right to reasonable inferences and draw attention to the difficulty of protecting people against assessments whose significance lies beyond the original information collected about them.⁵ Their work helps explain why a legal regime centred on collection may be insufficient when analytics creates an assessment that the person never supplied. The present article accepts that starting point without presenting it as a new discovery. It asks a narrower question about institutional responsibility after an inference has been created. If one organisation generates the assessment and another uses it, a challenge must reach both the evidence supporting the assessment and the decision influenced by it. A right framed only against the visible organisation may fail when that organisation cannot inspect the model, while a right framed only against the supplier may fail to change the adverse decision.

Selbst and Barocas distinguish the difficulty of describing a machine learning system from the difficulty of evaluating whether its decision rules are justified.⁶ This distinction prevents explanation from becoming an end in itself. A plain account of the variables used can help a person find a factual error, but it does not necessarily justify the selection of those variables or their use in a particular setting. Keller and Aplin, in their 2026 examination of AI public transparency, consider trade secret barriers and the role of external participation and institutional stewardship.⁷ Their analysis also cautions against assuming that broad exclusions from secrecy will solve every accountability problem. This article addresses individual contestation within a commercial vendor chain, rather than the broader public right to know. It extends the institutional question to the practical completion of a remedy, including the recipients who continue using an assessment after its original basis has been challenged.

The contribution is deliberately limited. The article does not claim that the protection of personal inferences or independent review of secrecy claims is a wholly new idea. It brings these issues together in a consumer-facing framework that links the profile, the commercial decision, the confidentiality claim and the outcome of a challenge. It distinguishes the supplier that can reconstruct a profile from the business that decides how to use it. The article's India-specific contribution is a legal and institutional proposal, not a claim that consumer attitudes or industry practices have been empirically measured. The argument

would be weakened by demanding disclosure of every algorithm or by assuming that European rules can be transferred to India without adaptation.

4. AI-GENERATED CONSUMER PROFILES AS A DISTINCT OBJECT OF PROTECTION

A personal inference can take the form of a score, a predicted attribute, a category or a narrative assessment. Its defining feature for this article is that it is attached to an identifiable person and is capable of influencing how that person is treated. A prediction need not be certain to matter. A tentative label can change the order in which an application is reviewed, the terms offered or the level of scrutiny applied. The distinction between a recorded fact and an inference therefore affects what correction can sensibly mean. A mistaken date can usually be replaced with the correct date. A prediction about future conduct may require a different response, such as correction of its inputs, reassessment of its relevance, disclosure of uncertainty or withdrawal from a use for which its reliability has not been demonstrated. An appropriate legal process should allow these possibilities without promising that every disputed opinion can be converted into an objectively true statement.

The Indian definition of personal data in section 2(t) refers to data about an individual identifiable by or in relation to such data.¹ A profile linked to an identified applicant presents a strong textual case for inclusion even when the applicant did not provide the profile herself. That is an interpretive argument, not a claim that the Act expressly names every form of inference. A model and its parameters require separate assessment. A general tool does not automatically become personal data merely because personal information was used in its development. Equally, a technical description does not settle whether identifiable information can be extracted or associated with people. The European Data Protection Board's Opinion 28/2024 treats the anonymity of AI models as a matter requiring assessment in the particular circumstances.⁸ Its relevance here is methodological: legal classification should follow evidence about the information and its use, rather than a blanket label attached to AI as a technology.

The proposed framework begins with the deployed personal output because it is usually the most direct point at which a person can identify a grievance. This does not imply that training and collection are unimportant. It means that a remedy should not depend on proving the entire development history of a system before an individual can ask why a particular assessment affected her. The organisation should be able to identify the assessment it relied upon, distinguish it from information supplied by the person and explain its practical role. If a profile has been updated repeatedly, the relevant record is the version used at the time of the challenged decision. A later and more accurate profile cannot explain an earlier refusal unless the relationship between the two is established. Version history is therefore relevant to accountability even when the legal complaint concerns a relatively ordinary factual error.

5. COMMERCIAL CONFIDENTIALITY AND CONSUMER ACCOUNTABILITY

Trade secret protection can serve a legitimate purpose in the development and supply of analytical systems.

Businesses may invest in methods, confidential datasets and technical processes whose disclosure would enable competitors to appropriate their work. Article 39 of the TRIPS Agreement links protection of undisclosed information to secrecy, commercial value because of secrecy and reasonable steps to maintain confidentiality.⁹ These conditions direct attention to the particular information claimed as secret. They do not justify treating every document produced inside a business as immune from scrutiny. In the present setting, the useful question is whether disclosure of specified information about an individual's assessment would reveal a protected commercial asset and, if so, whether a less extensive disclosure could meet the person's legitimate need. That enquiry is more precise than asking whether the company owns the score. Commercial control over a system does not by itself answer a question about legal duties arising from its use.

The European Trade Secrets Directive separately recognises legally required disclosure and provides mechanisms to preserve confidentiality in proceedings.¹⁰ The Law Commission of India's Report No. 289 recommends legislation on trade secrets and discusses the Indian legal position, alongside comparative approaches and exceptions.¹¹ The report is a law reform source and its appended proposal is not treated here as enacted legislation. It supports a careful distinction between protecting a commercial interest and creating an unrestricted power to withhold information relevant to a legal claim. In a vendor chain, the supplier's confidentiality agreement may explain why the deployer lacks access to certain evidence. It should not be allowed to determine, without further enquiry, whether the affected person has a valid challenge. A contract can distribute commercial obligations between the parties, but it cannot conclusively resolve the scope of an individual's statutory protection.

The ownership vocabulary is particularly unhelpful when applied to personal assessments. An analytics provider may have a legitimate claim in its methods while the person remains exposed to the consequences of a profile generated by those methods. The interests coexist. The proposed framework therefore does not require a transfer of the model or a general licence to exploit the supplier's technology. It requires access appropriate to the task of understanding and contesting the use of an assessment. The supplier should identify which part of the requested material carries the secrecy interest and explain the anticipated harm from disclosure. The deployer should identify why the assessment was relevant to its decision. Separating those responsibilities prevents the cost of developing the tool from becoming an assumed justification for every downstream use. It also protects suppliers against unnecessarily broad requests whose scope bears little relationship to the individual's grievance.

6. THE INDIAN LEGAL FRAMEWORK AND CONSUMER REDRESS

Section 11 of the Indian Act provides a right to obtain a summary of personal data and processing activities, information about specified recipients and other information as may be prescribed. Its consent related gateway, including the reference to section 7(a), is important.¹ The provision should not be described as an

unrestricted right available in identical form across every statutory basis for processing. Nor does its text contain the same express entitlement to meaningful information about automated logic found in the GDPR. A summary could reasonably identify a personal score that is being processed, but that does not necessarily settle what explanation of the score must accompany it. The argument for including deployed inferences in an access summary is therefore stronger than an argument that section 11 already requires disclosure of every model component. Conflating the two would obscure the precise point at which interpretation gives way to a proposal for reform. Section 8(3) addresses completeness, accuracy and consistency where personal data is likely to be used to make a decision affecting the Data Principal or disclosed to another Data Fiduciary. Section 12 addresses correction of inaccurate or misleading personal data, completion, updating and erasure within its statutory conditions.¹ These provisions offer an important starting point for examining a deployed profile. They nevertheless leave questions about evaluative predictions that are difficult to classify as simply accurate or inaccurate. A statistically plausible assessment may be unsuitable for an unrelated purpose, based on a comparison that does not fit the individual or presented with more certainty than the evidence supports. The proposed framework should therefore distinguish a request to correct factual information from a request to reconsider a prediction or suspend its use. The latter may require express authority beyond a duty to correct an inaccurate record. Legal precision requires that limitation to remain visible throughout the analysis.

The Act also addresses the relationship between a Data Fiduciary and processing carried out on its behalf, including the use of a processor under a valid contract. The designation of a supplier depends on its actual role in determining purpose and means, rather than the commercial label used in an agreement.¹ Rule 13 of the notified Rules adds an important prospective element for Significant Data Fiduciaries through assessment, audit and due diligence concerning technical measures, including specified uses of algorithmic software.³ The rule is relevant to supplier review, but it does not establish a general individual right to a model explanation. The Removal of Difficulties Order dated 5 October 2026 replaces the reference to periodic audit in section 10(2)(c)(ii) with data audit.¹² That clarification reinforces the statutory focus on data compliance. It does not establish a general review of the commercial merits of every decision made with AI. The analytical conclusion is limited: the enacted framework contains several foundations for documenting and reviewing personal inferences, while a complete procedure for consequential contestation across independent organisations remains a question for clarification and, in parts, legislative development. The proposed approach should supplement those foundations without enlarging administrative powers by assertion.

7. COMPARATIVE LESSONS FOR CONSUMER TRANSPARENCY

The GDPR expressly addresses automated decision making in Article 22 and information about its logic in

Article 15(1)(h), while Article 16 provides a right to rectification.¹³ These provisions operate within their own thresholds, conditions and exceptions. They should not be read as a universal right to explanations of every computer assisted decision. Their comparative value is that they identify the relationship between understanding the processing and challenging its effect. The Indian Act uses different language and institutional arrangements. The comparison therefore offers a design question rather than a ready answer: if a person has a legally recognised opportunity to challenge consequential processing, what information must be available to make that opportunity effective? A useful answer should identify the output actually used, its role in the decision and the information necessary to expose a material error, while recognising that the complete technical system may contain information unrelated to the individual's claim.

In *Dun and Bradstreet Austria*, decided on 27 February 2025, the Court of Justice interpreted meaningful information about automated logic as an intelligible account of the procedure and principles actually applied to obtain a particular result. Where a controller claimed that relevant material contained protected third-party information or trade secrets, the Court required the allegedly protected information to be supplied to the competent authority or court for a balancing assessment.¹⁴ This approach is significant because the business does not finally determine the boundaries of access merely by asserting confidentiality. It also distinguishes an explanation useful to the individual from material required by an expert reviewer. The judgment is not binding Indian law. Its narrower lesson for the present proposal is that secrecy disputes need an institutional decision maker with access to the evidence and authority to determine the appropriate form of disclosure.

The comparative approach still needs adaptation to the vendor chain. The organization using the score may hold the decision record but lack the supporting model evidence, while the supplier may possess that evidence but have no authority to reverse the decision. An Indian framework should expressly address that division. The aim should be an accountable process in which the body reviewing secrecy can also establish whether the individual received an adequate response and whether any correction reached the relevant users. A confidential review that validates a model in general would not answer a complaint that the wrong person's record was matched to an applicant. Likewise, a correction made by the supplier would not remedy a refusal if the deployer continued using an obsolete result. The legal process must be organized around the challenged assessment and its consequences, even when technical review requires broader access to the system.

8. A CONSUMER-FOCUSED RECORD OF THE PROFILE AND ITS USE

The first component of the proposal is a contemporaneous record connecting a consequential inference to the decision in which it was used. The record should identify the person or account to which the assessment was attached, the assessment itself, its date and version, the material input categories, the supplier and the organization responsible for the decision. It should also

state the purpose for which the assessment was used and the significance assigned to it. This is a proposal for a documentation standard, not a statement that the Act currently prescribes every field. The purpose is to prevent a grievance from becoming an enquiry into an indeterminate system. A company should be able to establish whether a profile was decisive, one factor among several or merely displayed without being acted upon. The appropriate response may differ in each situation, and the distinction should be supported by records created when the decision occurred.

The individual's account of that record should be clear enough to support a specific challenge. In a hypothetical screening process, the applicant should be told that an assessment treated a gap in her recorded employment as evidence relevant to continuity of service, if that was what actually happened. She could then explain that the record omitted a period of work. A generic statement that an AI model considered many factors would not enable the same challenge. Equally, the account should not promise that changing one factor will always change the result. Its wording must reflect the system actually used and any limits on reconstruction. If a supplier can only provide an approximate account of influential factors, that limitation should be disclosed rather than converted into a confident causal explanation. The standard should reward faithful information, including a candid statement of uncertainty, instead of demanding a polished narrative unsupported by the technical evidence.

The supplier's internal record would be more detailed than the account given to the individual. It should preserve information sufficient for a competent reviewer to examine the contested assessment, with retention proportionate to legal obligations and the likelihood of a dispute. The proposal does not justify indefinite storage of all personal inputs or every historical model. A system designed for reproducibility could instead preserve the relevant version identifier, input provenance and controlled evidence of the operation that produced the result. Where exact reconstruction is impossible, the organization should document that limitation and identify what evidence remains available. This creates an incentive to consider reviewability at procurement and design stages. A deployer selecting a supplier would need to ask whether the supplier can support legal scrutiny, rather than discovering only after a complaint that the necessary records were never retained.

9. REVIEWING CONFIDENTIALITY WITHOUT DISCLOSING THE ENTIRE MODEL

The second component is a reasoned assessment of confidentiality. A supplier seeking to withhold information should identify the material, explain the basis of the secrecy claim and describe the specific risk posed by the requested disclosure. The assessment should consider whether an intelligible summary, redaction or controlled expert access would answer the grievance with less commercial harm. These alternatives should not be treated as automatically adequate. A summary that conceals the relevant error fails its purpose, while an unnecessarily expansive disclosure may expose both commercial information and personal data concerning others. The standard is functional sufficiency: does the

information allow the particular dispute to be evaluated? Under the proposal, the supplier bears responsibility for justifying its secrecy claim because it possesses the relevant knowledge. The individual should not have to identify confidential technical details that she cannot know exist before obtaining any meaningful review.

An independent reviewer should be able to examine contested evidence under appropriate confidentiality controls. Access may require a secure environment, restrictions on copying, a defined review purpose and sanctions for misuse, depending on the material and the legal authority supporting the process. These measures are safeguards for a proposed procedure, not powers assumed to exist in every Indian forum. The reviewer should then provide a reasoned outcome explaining whether the individual received sufficient information, whether a material error was found and what remedial action is required. A confidential annex may preserve technical detail, but the person needs an account of the reasons sufficient to understand and contest the outcome. A process in which the reviewer merely assures the individual that the system is satisfactory would reproduce the original dependence on an inaccessible authority. Confidentiality should narrow the form of reasons only to the extent justified by a specific protected interest.

The reviewer must also distinguish commercial sensitivity from other grounds for limiting disclosure. The requested material may contain information about other people, security vulnerabilities or methods whose disclosure could facilitate manipulation. Each concern requires its own assessment. For example, a request about whether an applicant's record was confused with another record may be resolved by confirming the matching error without revealing the other person's details. A request about a suspect threshold may require more technical scrutiny. This approach avoids the assumption that either unrestricted access or complete secrecy is always appropriate. It also makes the review outcome auditable. A later appeal should be able to identify the interest protected, the alternative considered and the reason the chosen form of disclosure was sufficient. That discipline is particularly important when several suppliers make overlapping secrecy claims and the deployer cannot explain which claim prevented which part of the response.

10. CORRECTION ACROSS THE COMMERCIAL SUPPLY CHAIN

The third component is evidence that an accepted challenge has altered the relevant record and its uses. Correcting an input, recalculating an assessment, withdrawing a profile and reconsidering a decision are different actions. The remedy should identify which action is necessary and who can perform it. A supplier may be capable of issuing a replacement score while the deployer must decide whether to reopen an application. If the original score was shared with other organizations', its continued use becomes a separate concern. The proposed framework would require a record of the relevant recipients and a documented response to the correction, subject to a lawful and proportionate retention policy. This is a recommended coordination duty whose full scope would require legal authority. It should not be represented as an existing universal entitlement to recall

every inference under the Indian Act.

The individual should receive a completion statement that describes what changed, which relevant recipients were notified and whether any identified use remains unresolved. Such a statement would not guarantee that all historical effects can be reversed. A job may already have been filled or a time limited offer may have expired. The value of the statement is to distinguish a completed data correction from an outstanding dispute about the consequential decision. Where a business retains a historical assessment for a legal purpose, it should distinguish that archival record from a profile available for active use. A correction process should not silently delete evidence required to establish that the earlier decision was made on an incorrect basis. The proposed approach therefore combines correction with preservation of an appropriately restricted history. It seeks to stop continuing harm while retaining the evidence needed to evaluate what already occurred.

Responsibility should follow function. The organization making the consequential decision should answer for the role assigned to the inference and for any reconsideration within its control. The supplier should answer for the information it generated, its ability to support review and the execution of changes within its systems. If both independently determine purposes and means, their obligations cannot be reduced to a simple relationship between a fiduciary and a processor. Procurement contracts should specify support for access requests, preservation of relevant evidence, notification of accepted corrections and cooperation with authorized review. They should also state what happens when a supplier cannot provide the promised assistance. These provisions would make the proposal operationally credible without pretending that private contracts create rights against every recipient. The wider legal framework remains necessary to prevent a chain of individually reasonable contracts from leaving the affected person without an effective point of responsibility.

11. CONSUMER-FACING EXAMPLES AND PRACTICAL OBJECTIONS

The proposed framework can be tested against two hypothetical situations. In the first, a lender uses a supplier's assessment that incorrectly associates an applicant with another person's repayment history. The central issue is record matching. A useful response identifies the assessment and the matching basis, allows correction of the factual association and documents the supply of a revised assessment to the lender. Inspection of the full model would usually add little to that task. In the second, a recruitment service supplies a prediction about likely continuity of employment based on an incomplete work history. Correcting the history may resolve part of the problem, but the applicant may also dispute the relevance of continuity to the role. That question concerns the deployer's decision policy as well as the supplier's prediction. A remedy confined to deleting a field would not establish whether the application was reconsidered on a defensible basis.

These examples show why the scope of review should be tied to the disputed issue. They do not demonstrate that every predictive judgement is unlawful or that every

explanation can be produced at low cost. Smaller organizations may face genuine difficulty negotiating access from dominant suppliers. Suppliers may also face multiple requests whose cumulative effect reveals sensitive methods. The proposal addresses these concerns through materiality and controlled disclosure, but those safeguards need empirical evaluation. A pilot could examine the cost of maintaining the relevant records, the time needed to resolve disputes and the proportion of successful corrections that reach downstream users. Until such evidence exists, claims about reduced costs or improved accuracy would be premature. The doctrinal conclusion is narrower: a framework that leaves evidence and remedial authority disconnected has an identifiable structural weakness, even if the prevalence and magnitude of the resulting harm remain unmeasured.

There is also a risk that documentation becomes a substitute for substantive review. A supplier could provide an elaborate explanation that accurately describes an unsuitable practice. The deployer could record that it considered a profile without examining whether the profile should have been used. The proposed completion statement cannot by itself solve these problems. Its value depends on an independent process capable of asking the relevant legal and factual questions and on a decision maker with authority to require an appropriate response. Public reporting of aggregate complaint outcomes could reveal recurring problems without disclosing individual profiles or confidential methods. Such reporting should avoid creating another dataset from which vulnerable people can be identified. The framework therefore requires attention to both the privacy of the complainant and the quality of the decision under review. Better paperwork is useful only insofar as it enables a better determination of the grievance.

12. OPTIONS FOR INDIAN REFORM AND INSTITUTIONAL LIMITS

Several elements could be encouraged through implementation guidance without treating them as newly created rights. Organizations could identify consequential personal outputs in their data inventories, record recipients and include cooperation with lawful access and correction requests in supplier contracts. Significant Data Fiduciaries could examine inference related risks and supplier reviewability within the assessment and audit structure established by the Act and notified Rules.^{1,3} Guidance should explain its legal basis and avoid converting an interpretation into an unsupported mandate. Prescribing additional information under section 11(1)(c) may offer a route for more informative access, but the scope of that power must remain consistent with the parent statute. A general right to reconsider consequential automated assessments across every processing basis cannot safely be assumed to arise from a provision concerned with additional access information.

More substantial reform would require express legislative choices. Parliament could define a right to contest specified consequential personal inferences, establish the corresponding disclosure and cooperation duties and authorize an appropriate forum to review confidentiality claims. The reform should address the rights gateway, the role of independent suppliers, the conditions for

suspending disputed use and the relationship with other sectoral remedies. It should also provide a lawful basis for controlled expert access and procedures for protecting personal and commercial information obtained during review. The present article does not assume that the Data Protection Board already has a general power to reverse every service or employment decision. Its role would have to be established within the statutory scheme. Data correction, review of a commercial decision and compensation for an established injury are distinct remedial questions and may require coordination between different legal processes.

The limits of the proposal should guide its future development. A consequential threshold must be sufficiently clear to avoid disputes about every trivial recommendation while still covering assessments that materially shape opportunities. A purely monetary threshold could overlook serious effects on dignity or participation. A purely subjective threshold could make routine administration unpredictable. Sector specific standards may therefore be useful, provided they do not weaken a general minimum of meaningful access and correction. Further research should examine actual procurement contracts, the ability of suppliers to reconstruct assessments and the experience of people who attempt to contest profiles. Those enquiries would require legitimate access to data and appropriate research safeguards. They would test whether the legal allocation proposed here is workable. This article supplies a framework for that investigation, rather than claiming that doctrinal analysis has already answered the empirical questions.

13. CONCLUSION

Indian data protection law should address the gap between the business using an AI-generated consumer profile and the supplier capable of explaining how it was produced. Recognising a profile linked to an identifiable person as personal data is an important starting point, but meaningful protection also depends on whether the consumer can obtain useful information and an effective remedy. The article proposes a connected process: a record of the profile and its use, a specific and reviewable justification for withholding confidential information, and evidence that an accepted correction reaches relevant recipients. This approach protects the distinction between personal outputs and proprietary technical assets. It does not treat disclosure of an entire model as the normal measure of accountability.

The staged implementation of the Indian framework offers an opportunity to develop clearer practices while respecting the limits of existing law. Documentation, procurement terms and carefully grounded guidance may support better handling of consumer challenges. A comprehensive right to contest specified consequential profiles, together with enforceable cooperation and remedial duties across independent businesses, would require clear legal authority. European materials provide useful examples of reviewable confidentiality, but they do not remove the need for an Indian institutional design. The practical test is whether a consumer can identify a material problem, have it examined by an appropriate decision-maker and obtain a response that changes the relevant

processing or commercial decision. This is a conceptual legal proposal. Further empirical research should examine consumer experiences, business practices and the costs of implementing such a process.

REFERENCES

- 1 Government of India, The Digital Personal Data Protection Act, 2023, Act 22 of 2023, sections 2(i), 2(t), 7(a), 8(1) to 8(3), 10, 11, 12 and 13, <https://www.indiacode.nic.in/bitstream/123456789/22037/1/a2023-22.pdf> (accessed 10 October 2026)
- 2 Ministry of Electronics and Information Technology, Notification GSR 843(E), 13 November 2025, Gazette of India Extraordinary, commencement of the Digital Personal Data Protection Act, 2023, <https://www.meity.gov.in/static/uploads/2025/11/c56ceae6c383460ca69577428d36828b.pdf> (accessed 10 October 2026)
- 3 Ministry of Electronics and Information Technology, Digital Personal Data Protection Rules, 2025, GSR 846(E), 13 November 2025, rules 1, 13 and 14, <https://www.meity.gov.in/static/uploads/2025/11/53450e6e5dc0bfa85ebd78686cadad39.pdf> (accessed 10 October 2026)
- 4 Ministry of Electronics and Information Technology, Corrigenda to the Digital Personal Data Protection Rules, 2025, GSR 892(E), dated 10 December 2025, Gazette of India Extraordinary, 11 December 2025, <https://www.meity.gov.in/static/uploads/2025/12/3c7ebbae0e5456f493f486e6845df86b.pdf> (accessed 10 October 2026)
- 5 Wachter S and Mittelstadt B, A Right to Reasonable Inferences: Re Thinking Data Protection Law in the Age of Big Data and AI, *Columbia Business Law Review*, 2019(2) (2019) 494 to 620, DOI 10.7916/cblr.v2019i2.3424, <https://ora.ox.ac.uk/objects/uuid:d53f7b6a-981c-4f87-91bc-743067d10167> (accessed 10 October 2026)
- 6 Selbst A D and Barocas S, The Intuitive Appeal of Explainable Machines, *Fordham Law Review*, 87(3) (2018) 1085 to 1139, <https://ir.lawnet.fordham.edu/flr/vol87/iss3/11/> (accessed 10 October 2026)
- 7 Keller P and Aplin T, Reconciling Trade Secrets and AI Public Transparency, *Cambridge Law Journal*, 85(1) (2026) 151 to 185, DOI 10.1017/S0008197325101086, <https://kclpure.kcl.ac.uk/portal/en/publications/reconciling-trade-secrets-and-ai-public-transparency/> (accessed 10 October 2026)
- 8 European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models, adopted 17 December 2024, section 3.2, https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf (accessed 10 October 2026)
- 9 World Trade Organization, Agreement on Trade Related Aspects of Intellectual Property Rights, Article 39(1) and 39(2), https://www.wto.org/english/docs_e/legal_e/31bis_trips_04d_e.htm (accessed 10 October 2026)

- 10 European Parliament and Council, Directive (EU) 2016/943 of 8 June 2016, Articles 1(2), 2(1), 3(2) and 9, Official Journal of the European Union, L 157 (2016) 1 to 18, <https://eur-lex.europa.eu/eli/dir/2016/943/oj/eng> (accessed 10 October 2026)
- 11 Law Commission of India, Trade Secrets and Economic Espionage, Report No 289 (March 2024), covering recommendation and discussion of privacy in paragraphs 2.18 to 2.22, https://delhihighcourt.nic.in/files/2025-05/reports/289_lcir_part_1.pdf (accessed 10 October 2026)
- 12 Ministry of Electronics and Information Technology, Digital Personal Data Protection (Removal of Difficulties) Order, 2026, SO 5458(E), dated 5 October 2026, Gazette of India Extraordinary, 6 October 2026, paragraph 2(ii), <https://www.meity.gov.in/static/uploads/2026/10/daf495a7b5798a3dfa72a944b51fc176.pdf> (accessed 10 October 2026)
- 13 European Parliament and Council, Regulation (EU) 2016/679 of 27 April 2016, Articles 15(1)(h), 16 and 22, Official Journal of the European Union, L 119 (2016) 1 to 88, <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng> (accessed 10 October 2026)
- 14 Dun & Bradstreet Austria, Case C 203/22, ECLI:EU:C:2025:117, judgment of 27 February 2025, paragraphs 59 to 66 and 73 to 76 and operative part, https://juris.curia.europa.eu/juris/document_print.jsf?docid=295841&doclang=EN (accessed 10 October 2026)