



The Indian Journal for Research in Law and Management

Open Access Law Journal – Copyright © 2026

Editor-in-Chief – Dr. Muktai Deb Chavan; Publisher – Alden Vas; ISSN: 2583-9896

This is an Open Access article distributed under the terms of the Creative Commons Attribution-Non-Commercial-Share Alike 4.0 International (CC-BY-NC-SA 4.0) License, which permits unrestricted non-commercial use, distribution, and reproduction in any medium provided the original work is properly cited.

MACHINES MAY DETECT, OFFICERS MUST DECIDE: ALGORITHMIC ASSISTANCE AND THE LIMITS OF DELEGATION IN GST ADMINISTRATION

~ Harsh Jain, Mumal Singh and Apoorva Jamra

Abstract

India's Goods and Services Tax (GST) already runs on algorithms: returns are matched by machine, taxpayers are ranked by risk, and invoices are authenticated through a central portal. Within seven weeks in mid-2026, two High Courts quashed GST proceedings tainted by artificial intelligence (AI), and in doing so raised a question that tax law had not previously needed to answer: who is the decision-maker when a machine has done much of the thinking? This article argues that the answer lies in a familiar principle of administrative law rather than in a new body of 'AI law'. A statutory officer may use an algorithm as evidence or as a drafting aid, but may not allow it to stand in for the application of mind that the statute demands. The article distinguishes predictive analytics from generative drafting tools, shows that the recent decisions address only the second, and warns that a doctrine triggered by visible error will miss competent but unaccountable automation. It proposes an 'algorithm-removal test' and a framework of human-accountable algorithmic taxation to give the principle operational content.

Keywords: *Goods and Services Tax; artificial intelligence; non-delegation; application of mind; reasoned decisions; natural justice; algorithmic bias; tax administration*

I. Introduction

In July 2026 the Punjab and Haryana High Court quashed a GST notice because of something the department never intended anyone to read. Uploaded to the portal alongside the notice was an annexure of artificial intelligence (AI) suggestions on how the order might be strengthened,

with case law to add and rebuttals to anticipate.¹ Six weeks later the Gujarat High Court set aside a registration cancellation, and the appellate order that upheld it, after a probationary State Tax Officer conceded that the judgments cited in his order had been produced with AI and either did not exist or did not support what he said they did.²

It would be easy to read these decisions as a warning about careless officers and unreliable chatbots. That reading is too small. GST was designed as a technology-driven tax, and the administration has for years used computational methods to select returns, match invoices and rank taxpayers by risk. The cases therefore sit at the front edge of a much larger question: when public power is exercised through computational systems, what must remain human for the exercise to remain lawful?

This article argues that the answer is already contained in the rule against delegation of statutory discretion and in the requirement of application of mind. An algorithm may detect, rank and suggest; the statutory officer must decide, and must be able to show that the decision is the officer's own. The argument proceeds in three steps. It first separates two kinds of system that the present debate tends to run together, namely predictive analytics and generative drafting tools. It then examines what the courts have and have not decided, and tests the principle against scrutiny, reasons, natural justice and Article 14 of the Constitution. Finally, it proposes a practical framework, built around what I call the algorithm-removal test, for translating the principle into administrative practice.

Two contributions are claimed. The first is analytical: the recent cases concern generative tools, whereas the larger structural risk lies in predictive systems, and the legal response must differ accordingly. The second is critical: a rule that bites only when AI use is visible will protect taxpayers against the careless officer and leave them exposed to the careful one.

II. Algorithmic GST: Two Families of System

GST administration became algorithmic well before generative AI arrived. Section 61 of the Central Goods and Services Tax Act 2017 empowers the proper officer to scrutinise returns and to communicate discrepancies to the taxpayer, ordinarily in Form GST ASMT-10.³ In practice, the returns to be scrutinised are increasingly chosen at the centre. CBIC Instruction

¹ *M/s SRO India v State of Punjab and Another* (Punjab and Haryana HC, 7 July 2026).

² *Faiz Enterprise through Proprietor Mahetar Mahir Farukbhai v State Tax Officer, Unit-67, Surat* (Special Civil Application No 10123 of 2026, Gujarat HC, 20 August 2026), 2026:GUJHC:53490-DB.

³ Central Goods and Services Tax Act 2017, s 61; Central Goods and Services Tax Rules 2017, r 99.

No 02/2022-GST provides for selection on the basis of risk parameters, with the Directorate General of Analytics and Risk Management identifying GSTINs and passing the underlying data and revenue implications to field formations.⁴ E-invoicing deepens the process: each eligible invoice receives an Invoice Reference Number on registration, converting a commercial document into structured data that can be matched against returns and e-way bills.⁵ Practitioner commentary summarises these developments as risk profiling, pattern detection and process automation.⁶

Within this landscape it is useful to separate two families of system, because they fail in different ways and attract different legal responses. Predictive systems classify and rank. They ingest return, invoice and e-way bill data and output a flag, a score or a network of suspected linkages. Their characteristic failures are statistical: skewed or stale data, proxies that disadvantage particular groups of taxpayers, and scores whose basis cannot be explained. Generative systems, by contrast, produce text. They draft notices, summarise replies and suggest authorities. Their characteristic failures are textual: invented citations, plausible but unsupported reasoning, and the quiet replacement of the officer's own analysis with the model's. The distinction matters because the legal question changes. For predictive systems the question is evidential and procedural: what weight can a score bear, and what must the taxpayer be told? For generative systems the question is one of authorship: whose reasoning is on the page? Both questions reduce, in the end, to a single inquiry into who has decided. It is the second family that the courts have so far confronted.

III. What the Courts Have Decided

In *SRO India*, the petitioner challenged a pre-show-cause intimation in Form GST DRC-01A dated 2 February 2025.⁷ Annexed to the notice, and uploaded with it, was material generated by an AI tool, including suggestions on how the order could be improved and, reportedly, a

⁴ CBIC, Instruction No 02/2022-GST (22 March 2022).

⁵ Central Goods and Services Tax Rules 2017, r 48.

⁶ 'AI in GST: Powerful Tool, Unready Taxpayers: Risks of Algorithmic Judgement' (*TaxGuru*) <<https://taxguru.in/goods-and-service-tax/ai-gst-powerful-tool-unready-taxpayers-risks-algorithmic-judgement.html>> accessed 5 October 2026.

⁷ Central Goods and Services Tax Rules 2017, r 142(1A); 'P&H High Court Sets Aside GST Notice Which Contained AI Suggestions on How It Could Be Improved' (*Bar & Bench*, 21 July 2026) <<https://www.barandbench.com/news/ph-high-court-sets-aside-gst-notice-which-contained-ai-suggestions-on-how-it-could-be-improved>> accessed 5 October 2026.

prompt to make it ‘lethal’.⁸ The State told the court that these references had been uploaded inadvertently. The Division Bench was unmoved. Nothing before it showed that preparing and issuing a notice with an AI tool was permitted by the statute, and the notice had been issued primarily in reliance on such a tool. It was quashed together with the consequential proceedings.

Faiz Enterprise arose differently. The petitioner challenged a show-cause notice, an order cancelling its registration, the rejection of its revocation application and the appellate order.⁹ It pointed out that one of the judgments relied upon did not exist and that others carried incorrect citations or had no bearing on the issue. The court observed at an early hearing that the order appeared to have been passed by relying exclusively on AI-generated case law. The officer later filed an affidavit of unconditional apology, explaining that he was a probationary officer. The court quashed all four proceedings and directed fresh proceedings in which the defence would be considered and a reasoned order passed. It also took on record administrative instructions issued by the Additional Commissioner of State Tax on 18 August 2026 and directed that they be followed scrupulously, on pain of contempt. As reported, the instructions require officers to verify authorities against primary sources, to locate the original passage when AI paraphrases a judgment, to engage with authorities cited by the taxpayer, and to maintain human oversight, on the footing that AI may support reasoning but cannot replace the officer’s application of law to fact.¹⁰

Three features of these decisions deserve emphasis. First, neither court held that using AI is unlawful in itself. *Faiz Enterprise* is explicit that technology may assist; what failed was verification. *SRO India* is more troubling in its wording. The observation that the tool lacked ‘requisite sanction in law’ could, read literally, condemn every algorithmic aid not expressly authorised by statute.¹¹ That reading would imperil the risk-based selection on which scrutiny now depends. The better reading, and the only one consistent with the facts, is that the tool had

⁸ ‘AI Has No Statutory Authority: Punjab and Haryana High Court Nullifies GST Show Cause Notice Prepared Using Artificial Intelligence’ (*TaxO*, 18 July 2026) <<https://taxo.online/latest-news/18-07-2026-ai-has-no-statutory-authority-punjab-and-haryana-high-court-nullifies-gst-show-cause-notice-prepared-using-artificial-intelligence/>> accessed 5 October 2026.

⁹ ‘Gujarat High Court Quashes GST Orders After Officer Relies On AI-Generated Non-Existent Judgments’ (*Verdictum*, 27 August 2026) <<https://www.verdictum.in/gujarat-high-court/faiz-enterprise-v-state-tax-officer-2026gujhc53490-db-gst-orders-quashed-ai-generated-judgments-1620726>> accessed 5 October 2026.

¹⁰ ‘Gujarat High Court Flags Use of AI-Generated Case Laws in Tax Orders, Warns Tax Officers of Contempt’ (*Bar & Bench*, 27 August 2026) <<https://www.barandbench.com/news/litigation/gujarat-high-court-flags-use-of-ai-generated-case-laws-in-tax-orders-warns-tax-officers-of-contempt>> accessed 5 October 2026.

¹¹ *SRO India* (n 1).

displaced the officer's own consideration of the matter, so that the notice was not the product of the application of mind that the statute presupposes.

Secondly, both cases came to light through accident or visible error. A stray annexure and a fictitious judgment are the legal equivalent of fingerprints left at the scene. An officer who used a better tool, verified its output and uploaded only the polished result would leave no trace at all. A doctrine that depends on discoverable flaws therefore regulates carelessness rather than delegation, and it cannot tell a court whether a facially sound order reflects independent judgment or competent automation.

Thirdly, the remedy is modest. In each case the department may begin again. Quashing vindicates the principle, but it does little to alter the pressures that produce overreliance on a tool: heavy caseloads, limitation deadlines and the expectation that demands will be raised. The law needs a standard that officers can apply before the order is passed, and not merely one that courts apply afterwards.

IV. Detection Is Not Determination

The rule that emerges can be put simply. An algorithmic output is information, and information is not a finding. Nowhere is the distinction clearer than in scrutiny. A discrepancy between GSTR-1 and GSTR-3B may reflect an amendment, a credit note, a timing difference or a supplier's default as easily as under-reporting. A sudden rise in turnover may be evasion, or it may be a good year. The machine can say that a transaction merits attention; it cannot say that the taxpayer is liable. That is why section 61 frames scrutiny as the communication of discrepancies to the taxpayer for explanation, and why the demand provisions require the officer to determine tax actually unpaid or wrongly credited and, where the allegation is fraud, wilful misstatement or suppression, to establish it.¹²

Risk scoring sharpens the point. A score expresses a statistical association between a taxpayer's attributes and past non-compliance. A new registration, a high-value interstate supplier or a business in a sector with a history of default may attract a high score without having done anything wrong. The danger is circular: closer scrutiny of high-risk taxpayers uncovers more discrepancies, which feed back into the model as confirmation. Scholars of the 'scored society'

¹² Central Goods and Services Tax Act 2017, ss 61, 73, 74 and 74A.

have long warned that such loops entrench the very patterns they claim to detect,¹³ and the OECD identifies skewed data, weak explainability and improper targeting among the principal risks of AI in tax administration.¹⁴ A score may legitimately determine where the administration looks. It cannot itself establish that anything has been found.

Data quality compounds the problem. GST liability and input tax credit depend on information supplied by other taxpayers, and an honest recipient may appear suspicious because a supplier failed to file or reported incorrectly. A machine-learning model trained on such records does not lie; it learns accurately from inaccurate inputs and may present the result with misleading confidence. The evidential reliability of an algorithmic output can never exceed that of the data beneath it. It follows that the taxpayer must have a genuine opportunity to challenge the data, and not merely the conclusion drawn from it.

V. Delegation and Application of Mind

The legal foundation for these concerns is the ordinary rule that a statutory power must be exercised by the authority on whom it is conferred. The CGST Act vests scrutiny, assessment and adjudication in the ‘proper officer’.¹⁵ Administrative law has long distinguished between delegating mechanical tasks, which is unobjectionable, and abdicating judgment, which is not. Where a statute makes the officer’s opinion the foundation of action, the opinion must be genuinely formed on relevant material, and the Supreme Court has been prepared to examine whether it was.¹⁶ A show-cause notice, moreover, must not prejudge the outcome; it invites an explanation from an authority that has yet to decide. An annexure instructing a tool to make the order ‘lethal’ sits uneasily with that requirement.¹⁷

The difficulty is that application of mind is an internal state, and the law can test it only through external evidence. Here a human-in-the-loop requirement offers less comfort than it appears to. Research on automation bias shows that people reviewing machine recommendations tend to accept them and to search less for contrary evidence, particularly under pressure of workload.¹⁸ More recent work argues that policies requiring human oversight of government

¹³ Danielle Keats Citron and Frank Pasquale, ‘The Scored Society: Due Process for Automated Predictions’ (2014) 89 Wash L Rev 1.

¹⁴ OECD, *Governing with Artificial Intelligence* (OECD Publishing 2025), chapter on tax administration.

¹⁵ Central Goods and Services Tax Act 2017, s 2(91).

¹⁶ *Barium Chemicals Ltd v Company Law Board* AIR 1967 SC 295.

¹⁷ *Oryx Fisheries Pvt Ltd v Union of India* (2010) 13 SCC 427.

¹⁸ Linda J Skitka, Kathleen L Mosier and Mark Burdick, ‘Does Automation Bias Decision-Making?’ (1999) 51 Int J Hum-Comput Stud 991.

algorithms often fail because they assume capacities that overburdened reviewers do not have.¹⁹ A signature at the foot of a machine-shaped order therefore proves little. What is needed is a way of testing whether the decision survives without the machine.

I propose to call this the algorithm-removal test. Delete the algorithmic output from the file, whether a risk score, a discrepancy report or a draft, and ask whether the order still rests on identifiable primary material and on a chain of reasoning that the officer can articulate unaided. The test has two limbs. The evidential limb asks whether the facts relied upon can be traced to underlying records, such as invoices, returns and statements, rather than to the tool's characterisation of them. The reasoning limb asks whether each step from fact to statutory conclusion appears in the officer's own analysis, including the treatment of the taxpayer's reply. If either limb fails, the output has become the decision. The test does not ask whether AI was used, which is the wrong question, and it does not require officers to avoid tools. It asks whether, in the tool's absence, the officer could still defend the order.

VI. Reasons, Natural Justice and the Black Box

The Supreme Court has treated reasons as the link between material and conclusion, without which judicial review cannot function.²⁰ It has also insisted that an order must stand on the reasons it records and cannot be supplemented by explanations offered later.²¹ Where civil consequences follow, the requirements of a fair hearing apply whether or not the statute spells them out.²²

Algorithmic outputs strain these requirements because a probability is not a ground. To tell a taxpayer that a model rates it as high risk explains why the department looked; it does not explain why the taxpayer owes anything. The taxpayer is entitled to know which transactions and records are relied upon, what discrepancy is alleged, which provision is said to be breached, how the sum was computed and why its own explanation was rejected. The distinction between a selection explanation and a liability explanation is, I suggest, the crux. The former may legitimately be thin, because an administration need not publish the criteria by which it chooses whom to examine. The latter must be complete.

¹⁹ Ben Green, 'The Flaws of Policies Requiring Human Oversight of Government Algorithms' (2022) 45 Computer Law & Security Review 105681.

²⁰ *Kranti Associates Pvt Ltd v Masood Ahmed Khan* (2010) 9 SCC 496.

²¹ *Mohinder Singh Gill v Chief Election Commissioner* (1978) 1 SCC 405.

²² *Canara Bank v Debasis Das* (2003) 4 SCC 557.

Nor does the law require disclosure of source code. India has no statutory right to an explanation of automated decisions in the GST context, and the Digital Personal Data Protection Act 2023 contains no provision comparable to article 22 of the European Union's General Data Protection Regulation.²³ But the general law of fairness supplies much of what is needed. A reasoned account of the factual and legal basis of the decision, together with the material relied upon, allows the taxpayer to respond and the court to review without exposing proprietary models. Scholarship on algorithm-assisted public decisions has likewise stressed intelligible reasons and meaningful human intervention as the conditions of legitimacy, rather than transparency of code for its own sake.²⁴

VII. Article 14 and Algorithmic Classification

The use of a model to select taxpayers is State action and is subject to Article 14. A classification must rest on an intelligible differentia having a rational relation to the object of the law,²⁵ and State action must not be arbitrary.²⁶ The State cannot answer a complaint of discriminatory scrutiny by saying that the computer selected the taxpayer.

Risk-based selection is plainly rational in principle, since enforcement resources are finite. But a model that relies on proxies such as sector, location, age of registration or size may in practice burden some classes of taxpayer disproportionately, and the bias will not be visible on the face of any single notice. Human prejudice leaves traces in conduct; algorithmic bias hides in variables and weights. The State should therefore be able, if challenged, to show that selection criteria have been validated, that false-positive rates have been tested across relevant categories of taxpayer, and that systematic error can be corrected. I do not suggest that every selection decision invites constitutional litigation, which would defeat the purpose of risk-based administration. The point is narrower: a taxpayer who makes a credible showing of systematic disadvantage should be met with evidence of validation and not with a claim of technical neutrality.

²³ Digital Personal Data Protection Act 2023; Regulation (EU) 2016/679 [2016] OJ L119/1, art 22.

²⁴ Marion Oswald, 'Algorithm-Assisted Decision-Making in the Public Sector: Framing the Issues Using Administrative Law Rules Governing Discretionary Power' (2018) 376 Phil Trans R Soc A 20170359.

²⁵ *State of West Bengal v Anwar Ali Sarkar* AIR 1952 SC 75.

²⁶ *E P Royappa v State of Tamil Nadu* (1974) 4 SCC 3; *Maneka Gandhi v Union of India* (1978) 1 SCC 248.

VIII. Comparative Perspectives

Other jurisdictions have confronted the problem and offer useful, if imperfect, guidance. In *SCHUFA*, the Court of Justice of the European Union held that the automated generation of a credit probability score is itself an automated decision where a third party draws strongly on it in deciding whether to grant credit.²⁷ The reasoning translates well: where a risk score drives what an officer does, the score cannot be dismissed as a mere preliminary step. In the Netherlands, the Hague District Court struck down the SyRI legislation, which permitted data-linkage risk indication for welfare fraud, in part because the system's workings were too opaque to justify the interference with privacy.²⁸ And in *Loomis*, the Wisconsin Supreme Court allowed a risk assessment tool to be used at sentencing, but only with cautions about its limits and on the footing that it must not be determinative.²⁹

These authorities arise in different settings, but they share a theme. Courts tolerate algorithmic assistance and resist algorithmic determination, and they expect the State to be able to explain what the tool did. India need not wait for legislation to adopt the same stance, since the administrative-law principles discussed above already furnish a basis for it.

IX. Objections and a Framework

A fair objection concerns scale. GST administration handles volumes that no officer corps could review manually, and the system already produces automated outputs, such as system-generated notices to non-filers.³⁰ If every algorithmic step demanded human adjudication, the system would collapse. The answer is that the principle defended here is a gradient and not an absolute. Where an automated step is arithmetic, non-discretionary, visible to the taxpayer and easily corrected, a human signature adds little. As the step moves towards contested findings of fact, the characterisation of a transaction, an allegation of fraud or a penalty, the case for personal judgment strengthens, because these are the points at which the taxpayer's explanation matters and the cost of error is highest. Proportionality, in other words, governs how much humanity is required.

Four commitments would give the principle practical content. First, every system should be classified by function as detection, recommendation or determination. Detection and

²⁷ Case C-634/21 *OQ v Land Hessen (SCHUFA Holding (Scoring))* EU:C:2023:957.

²⁸ *NJCM c.s. v The Netherlands* (District Court of The Hague, 5 February 2020) ECLI:NL:RBDHA:2020:1878.

²⁹ *State v Loomis* 881 NW 2d 749 (Wis 2016).

³⁰ Central Goods and Services Tax Act 2017, s 46; Central Goods and Services Tax Rules 2017, r 68.

recommendation are permissible with safeguards; determination of liability, penalty or disputed fact should be presumptively impermissible where it displaces the statutory officer. Secondly, every notice or order materially informed by an algorithmic output should record the officer's reasoning in a form that passes the removal test, replacing formulaic certification of independent review with a record of what was verified against primary material. Thirdly, the administration should maintain an audit trail of data sources, extraction dates, model versions and the output relied upon, so that appellate authorities can reconstruct how a decision was reached,³¹ and should disclose algorithmic involvement to the taxpayer where it materially influenced the action. Fourthly, selection systems should be tested periodically for accuracy and disparate impact, with aggregate results reported.

The Gujarat instructions show that rules of this kind can be framed quickly, but they bind one State's officers and exist because a court required them. CBIC and the GST Council are better placed to issue uniform, binding guidance. None of this is hostile to technology. AI can also reduce compliance costs for taxpayers, for instance by assisting with classification and reconciliation before filing, and an administration that deploys it in that way earns the trust on which voluntary compliance depends.

X. Conclusion

The 2026 decisions are best understood as the application of old law to new tools, not as the creation of new law. A statutory officer must apply their own mind to the facts, must give reasons that the taxpayer can understand and contest, and must exercise discretion personally. None of that changes because a model has drafted the text or ranked the taxpayer. Yet the decisions also reveal the limits of a doctrine that detects delegation only when it is clumsy. The more consequential question for GST is not whether officers will be caught using AI badly, but whether the system can be designed so that using it well still leaves the officer as the author of the decision. The algorithm-removal test and the framework of human-accountable algorithmic taxation offered here are intended as steps towards that design. The governing proposition is plain: machines may detect, and officers must decide. The task for the law is to ensure that the second half of that sentence remains true in practice and not merely in form.

³¹ Danielle Keats Citron, 'Technological Due Process' (2008) 85 Wash U L Rev 1249; Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 Geo LJ 1147.