


COVER STORY



*Technology may flag the case.
People must decide it.*

DIRECT TAX · TECHNOLOGY & GOVERNANCE

Digital Tax Administration: Balancing Revenue Powers and Taxpayer Rights

As tax authorities turn to AI, predictive analytics and automated notices, a new friction emerges between enforcement efficiency and the taxpayer's right to due process. Here is a framework to reconcile the two.

By Sri Geetha Thiyagarajan, CA Final

IN THIS ISSUE

02

The rise of AI in direct tax administration

03

The asymmetric axis: when algorithms pressurise rights

05

A balanced paradigm: HITL, XAI and a Digital Ombudsperson

06

Comparative matrix & conclusion

ABSTRACT

This paper explores the structural paradigm shift in direct tax administration brought about by automated data processing, predictive analytics and artificial intelligence (AI). As sovereign tax authorities deploy advanced algorithmic architectures to maximise revenue intelligence, detect evasion and enforce compliance, a critical structural friction emerges between enhanced state enforcement capabilities and established constitutional taxpayer rights. Examining international algorithmic enforcement tools, the paper analyses where automated administration can inadvertently compromise procedural due process, transparency and the right to be heard – and proposes a balanced framework anchored in “Human-in-the-Loop” architecture, algorithmic explainability and statutory safeguards, to align fiscal efficiency with democratic equity.

KEYWORDS Digital Tax Administration · Artificial Intelligence · Revenue Enforcement · Taxpayer Rights · Procedural Due Process · Risk Profiling · Algorithmic Transparency · Black Box Analytics · Human-in-the-Loop · Explainable AI · Digital Ombudsperson · Fiscal Governance

01 Introduction

Traditionally, the balance between the State and the taxpayer was maintained through human agency: physical auditors reviewed paper ledgers, face-to-face assessments allowed for contextual explanations, and human discretion served as both a mechanism for enforcement and a buffer against arbitrary state action. However, the rapid digitisation of the global economy has generated data volumes that render manual oversight obsolete.

In response, revenue administrations worldwide have embarked on sweeping digital transformations. AI,

machine-learning models and Big Data analytics have redefined fiscal sovereignty: tax administration has evolved from a reactive, audit-based model into a proactive, real-time predictive enterprise. While these technologies eliminate tax leakage, suppress shadow economies and lower compliance friction for honest taxpayers, they raise a distinct challenge – *how do we preserve the equilibrium of the social contract when human discretion is replaced by automated algorithms?*

02 The technological architecture: AI in direct tax administration



Figure 1. The three core functional modalities through which AI operates in direct tax regimes.

To understand the legal challenges, one must first understand the mechanics of contemporary fiscal systems. Modern direct tax platforms use advanced neural networks to parse vast lakes of unstructured data. They execute automatic cross-verification, reconciling a taxpayer's self-declared income against actual consumption patterns, luxury asset acquisitions, foreign remittances and corporate cross-holdings. Processing data at scales impossible for human

teams, machine-learning models recognise complex patterns of transactional structuring that indicate systemic evasion or aggressive base erosion.

Predictive models also forecast compliance behaviour. By classifying taxpayers into granular behavioural clusters, the system decides which profiles receive automated text notifications, which receive formal electronic inquiries, and which warrant comprehensive forensic investigation.

03 The asymmetric axis: how algorithmic enforcement pressurises taxpayer rights

As state revenue engines gain unprecedented digital capabilities, the traditional legal safeguards designed to protect citizens face systemic strain. “The Asymmetric Axis” describes a structural imbalance of power: on one side, the State — armed with computational power, hidden algorithms and big data; on the other, an individual taxpayer who does not know *how* the machine flag was generated, yet bears the entire burden of proving it wrong.




Figure 2. Machine-driven conclusions can create financial and legal consequences before a human official ever reviews the case.

3.1

The dilution of procedural due process and the right to be heard

The right to be heard is a core pillar of natural justice: no person should be penalised without a fair chance to state their case. Yet when automated tax systems instantly flag an accounting mismatch, they often bypass human dialogue and immediately generate a standardised penalty notice. Online portals frequently restrict a taxpayer's reply to rigid dropdown menus or tiny text boxes, cutting off their ability to explain complex business context.

Illustration. An automated system flags a mismatched transaction from three years ago and claims ₹5,00,000 in unpaid taxes plus penalties. Instead of a formal notice and a scheduled hearing where the taxpayer's chartered accountant can present tax receipts, the authority proceeds through a fast-tracked, automated digital channel.

3.2

Information asymmetry and the opaque 'black box'

The “black box” problem is a transparency gap: an AI system can ingest millions of data points and penalise a taxpayer without revealing its internal logic. The department holds all the algorithmic insight; the citizen is left in an information vacuum, unable to pinpoint what triggered the decision.

When an AI flags a taxpayer on correlations across third-party data, it leaves no paper trail. The software simply issues an adjustment notice — “You owe Rs. 5,000.” If it does not specify which transaction or bank account caused the flag, the taxpayer is left fighting a ghost. For a tax system to be fair, a citizen must be shown not just *that* they are wrong, but exactly *how* the computer concluded they were wrong.

3.3

Algorithmic bias, data corruption and the presumption of guilt

In direct tax administration, a “presumption of guilt” arises when an AI system encounters a data mismatch and automatically assumes the taxpayer is hiding income — flipping the traditional burden of proof onto the citizen.

The problem is amplified by **data corruption**: when a bank makes a clerical error in a financial upload, the AI naturally trusts the third-party digital data over the citizen's accurate records. **Algorithmic bias** creeps in because models are trained on historical audit data; if past officers routinely targeted particular professions, the AI codifies that bias and keeps flagging the same cohorts.

Because the software works on historical predictive data that profiles independent freelancers as “high-risk”, it can ignore a consultant's genuine bank statements altogether. The innocent taxpayer is forced to run from pillar to post to prove their innocence — a reminder of how data errors and flawed training can turn a citizen into a suspect overnight.

“Predictive neural networks operate as black boxes — their internal processes are often opaque to the very people they impact.”

CASE IN POINT

One extra zero: a freelance consultant in Bengaluru

ACTUAL FEE EARNED

₹15,000

A legitimate payment for a small project.



REPORTED BY CLIENT

₹1,50,000

The client's accountant keys in an extra zero.



AI VERDICT

Income concealed

The mismatch is flagged and guilt is presumed.



CONSEQUENCE

₹30,000

Routine refund frozen, and a heavy penalty demanded.

04 The balanced paradigm: structural alternatives for equitable governance

Digital tax administration provides immense public benefits. It reduces corruption by minimising arbitrary human contact, dramatically lowers enforcement costs, stabilises public revenue and simplifies reporting for millions of citizens.

The core objective, therefore, is not to dismantle digital tools but to implement systemic guardrails that align algorithmic efficiency with the rule of law. Three structural alternatives, set out overleaf, can achieve that equitable balance.

**Human-in-the-Loop**

AI advises; a qualified officer decides and signs off.

**Explainable AI**

Every notice itemises the data behind the flag.

**Digital Ombudsperson**

Pauses enforcement and corrects data, fast.

At a glance. The three safeguards are detailed overleaf.

SECTION 04 · THREE SAFEGUARDS

Efficiency with accountability

Each safeguard tackles one face of the asymmetric axis: who decides, what the taxpayer is told, and how fast an error can be undone.



HUMAN JUDGMENT STAYS AT THE CENTRE

4.1 · WHO DECIDES



Implement 'Human-in-the-Loop' (HITL) validation

The most vital protection against automated error is ensuring that AI never holds final, unreviewed adjudicative authority over a citizen's legal liabilities. Algorithms must function strictly as diagnostic and advisory systems, while final determinations stay anchored in human oversight. When a model identifies a high probability of non-compliance, the system cannot automatically execute assessments, freeze assets or issue penalty orders. The output is routed to a qualified tax officer or oversight panel, who independently reviews the findings, evaluates the underlying data, checks the business context and manually signs off on any formal notice — preserving human accountability within automated processes.

4.2 · WHAT THE TAXPAYER IS TOLD



Mandate explainability (XAI) and legible notices

The "black box" of fiscal AI must give way to Explainable AI. Where an agency uses a predictive model to adjust a liability or select a profile for audit, the software must output a clear, human-legible explanation of its logic alongside its decision. An automated notice must break down the decision path. Itemised data points let the taxpayer identify the issue immediately, correct third-party errors or provide targeted documentation — restoring transparency to the process.

EXPLAINABLE NOTICE	SAMPLE
	"Your profile has been flagged because a transaction of Rs. 50,000 recorded by Bank X on Date Y does not correspond with the gross revenues reported under your return." Each highlighted field is a specific, checkable data point.

4.3 · HOW FAST AN ERROR IS UNDONE



Create a Digital Ombudsperson and fast-track correction

Digital errors propagate across interconnected state systems instantly, while traditional appeals take months or years. Governments should establish a dedicated statutory office — the **Taxpayer Digital Rights Ombudsperson** — empowered to receive immediate complaints about systemic glitches, algorithmic misclassifications and flawed data-matching. With specialised technical and legal expertise, it would have the authority to pause automated enforcement, correct corrupted data entries and give swift relief to citizens caught in algorithmic loops.

05 Comparative matrix: traditional, fully automated and balanced pathways

ADMINISTRATIVE VECTOR	TRADITIONAL FRAMEWORK	UNMITIGATED AUTOMATED AI	BALANCED APPROACH (PROPOSED)
Audit selection	Manual sampling and human discretion; prone to inconsistency and low coverage.	Opaque “black box” predictive models; prone to bias and systemic feedback loops.	Explainable AI risk scoring combined with human supervisory panel approvals.
Procedural due process	In-person hearings and full narrative arguments; high transaction costs.	Standardised e-notices with rigid binary inputs; little room for contextual defence.	Dynamic portals supporting structured data alongside rich contextual documentation.
Error correction	Formal appeals and long-term litigation; slow resolution.	Instantaneous automated enforcement; workflows hard to interrupt.	Independent Digital Rights Ombudsperson with rapid data-correction authority.

— Conclusion

The digitisation of tax administration is an irreversible and broadly beneficial evolution. AI gives states the capacity to build efficient, leak-proof revenue frameworks — and this can coexist with constitutional guarantees. Tech efficiency and taxpayer rights are not mutually exclusive; they are two sides of a well-functioning democratic state.

By integrating human-in-the-loop review, algorithmic transparency and agile dispute resolution, fiscal systems can remain both effective and equitable. The future does not require choosing between state efficiency and citizen rights: true progress lies in automated architectures that respect due process, growing more fair, transparent and accountable as they grow more powerful.

REFERENCES

- Income Tax Department, Government of India. *Faceless Assessment Scheme Guidelines and Implementation Framework*.
- Bird, R. (2020). *Taxation in the Digital Era: The Balancing Act Between Efficiency and Rights*. *Journal of Tax Administration*, 14(2), 112–135.



ABOUT THE AUTHOR

Sri Geetha Thiyagarajan

CA Final

SRO NO.

SRO0787496

EMAIL

srigeethaselvam2034@gmail.com

KEY TAKEAWAYS

- AI should advise; a human officer should decide.
- Every automated notice should show its working.
- Algorithmic errors need a fast, independent fix.

CONTACT

Old No 17B, New No 22/16A1,
Gandhi Road, Salem – 636007
shankar@teameyrie.com