

INTERNATIONAL JOURNAL FOR LEGAL RESEARCH AND ANALYSIS



Open Access, Refereed Journal Multi Disciplinary
Peer Reviewed

www.ijlra.com

DISCLAIMER

No part of this publication may be reproduced, stored, transmitted, or distributed in any form or by any means, whether electronic, mechanical, photocopying, recording, or otherwise, without prior written permission of the Managing Editor of the *International Journal for Legal Research & Analysis (IJLRA)*.

The views, opinions, interpretations, and conclusions expressed in the articles published in this journal are solely those of the respective authors. They do not necessarily reflect the views of the Editorial Board, Editors, Reviewers, Advisors, or the Publisher of IJLRA.

Although every reasonable effort has been made to ensure the accuracy, authenticity, and proper citation of the content published in this journal, neither the Editorial Board nor IJLRA shall be held liable or responsible, in any manner whatsoever, for any loss, damage, or consequence arising from the use, reliance upon, or interpretation of the information contained in this publication.

The content published herein is intended solely for academic and informational purposes and shall not be construed as legal advice or professional opinion.

**Copyright © International Journal for Legal Research & Analysis.
All rights reserved.**

ABOUT US

The *International Journal for Legal Research & Analysis (IJLRA)* (ISSN: 2582-6433) is a peer-reviewed, academic, online journal published on a monthly basis. The journal aims to provide a comprehensive and interactive platform for the publication of original and high-quality legal research.

IJLRA publishes Short Articles, Long Articles, Research Papers, Case Comments, Book Reviews, Essays, and interdisciplinary studies in the field of law and allied disciplines. The journal seeks to promote critical analysis and informed discourse on contemporary legal, social, and policy issues.

The primary objective of IJLRA is to enhance academic engagement and scholarly dialogue among law students, researchers, academicians, legal professionals, and members of the Bar and Bench. The journal endeavours to establish itself as a credible and widely cited academic publication through the publication of original, well-researched, and analytically sound contributions.

IJLRA welcomes submissions from all branches of law, provided the work is original, unpublished, and submitted in accordance with the prescribed submission guidelines. All manuscripts are subject to a rigorous peer-review process to ensure academic quality, originality, and relevance.

Through its publications, the *International Journal for Legal Research & Analysis* aspires to contribute meaningfully to legal scholarship and the development of law as an instrument of justice and social progress.

PUBLICATION ETHICS, COPYRIGHT & AUTHOR RESPONSIBILITY STATEMENT

The *International Journal for Legal Research and Analysis (IJLRA)* is committed to upholding the highest standards of publication ethics and academic integrity. All manuscripts submitted to the journal must be original, unpublished, and free from plagiarism, data fabrication, falsification, or any form of unethical research or publication practice. Authors are solely responsible for the accuracy, originality, legality, and ethical compliance of their work and must ensure that all sources are properly cited and that necessary permissions for any third-party copyrighted material have been duly obtained prior to submission. Copyright in all published articles vests with IJLRA, unless otherwise expressly stated, and authors grant the journal the irrevocable right to publish, reproduce, distribute, and archive their work in print and electronic formats. The views and opinions expressed in the articles are those of the authors alone and do not reflect the views of the Editors, Editorial Board, Reviewers, or Publisher. IJLRA shall not be liable for any loss, damage, claim, or legal consequence arising from the use, reliance upon, or interpretation of the content published. By submitting a manuscript, the author(s) agree to fully indemnify and hold harmless the journal, its Editor-in-Chief, Editors, Editorial Board, Reviewers, Advisors, Publisher, and Management against any claims, liabilities, or legal proceedings arising out of plagiarism, copyright infringement, defamation, breach of confidentiality, or violation of third-party rights. The journal reserves the absolute right to reject, withdraw, retract, or remove any manuscript or published article in case of ethical or legal violations, without incurring any liability.

USING ARTIFICIAL INTELLIGENCE TO VERIFY GST INVOICE AUTHENTICITY: TACKLING QR CODE FRAUD IN E-INVOICING

AUTHORED BY - ANKUL PRAJAPATI

CO-AUTHOR'S - SHUBHAM SHARMA & ARYAN PUNDIR

ABSTRACT

India's transition to a digitised tax ecosystem through GST and e-invoicing has strengthened tax administration and reduced reliance on manual verification. However, the existing e-invoicing framework contains a significant vulnerability while a digitally signed QR code authenticates the metadata embedded within it, it does not establish that the visual contents of the invoice correspond to that metadata. This "analog gap" enables manipulation of invoice amounts, GSTINs, and other fields while retaining a valid QR code, posing particular risks for Micro, Small and Medium Enterprises (MSMEs) that frequently rely on fragmented billing systems and digitally circulated invoices.

This paper examines the legal and technological implications of this verification gap, particularly in light of AI-assisted invoice forgery and QR-based phishing ("quishing"). It argues that reliance on QR-code authentication alone may expose genuine recipients of manipulated invoices to tax, penalty, and prosecution consequences despite their reliance on officially available verification mechanisms.

To address this problem, the paper proposes "GST Administration 3.0", an AI-enabled layered verification framework that supplements cryptographic authentication with Optical Character Recognition (OCR), semantic comparison, visual forensic analysis, and deep-learning-based anomaly detection. The proposed system would compare invoice contents with QR-encoded metadata, identify potential manipulation, generate risk classifications, and provide explainable verification reports. These records could also assist taxpayers in demonstrating due diligence and enable authorities to distinguish fraudulent actors from bona fide recipients. The paper further examines the statutory and constitutional implications of AI-enabled GST verification and proposes targeted reforms to GSTN verification mechanisms, CBIC enforcement guidance, and the e-invoicing framework. It concludes that AI-assisted verification can bridge the gap between cryptographic authentication and document integrity

while strengthening revenue protection and reducing disproportionate compliance risks for MSMEs.

Keywords: *GST, Tax Evasion, Digital Tax Reform, E-Invoicing, QR Code Fraud, MSME, AI Verification.*

INTRODUCTION

In 2023, the Directorate General of GST Intelligence (DGGI) uncovered a sophisticated fraud network comprising 23 bogus companies that systematically issued forged invoices to claim over ₹3,000 crore in fraudulent Input Tax Credit (ITC), marking one of the largest such scams in India's GST regime.¹ What rendered this case particularly alarming was not merely its unprecedented scale, but the simplicity of its execution: fraudsters modified invoice amounts, GSTIN details, and item descriptions on legitimate documents while retaining valid QR codes generated by the official Invoice Registration Portal (IRP). Downstream MSMEs, scanning these QR codes through the National Informatics Centre's (NIC) E-Invoice QR Code Verifier and receiving validation, claimed ITC in good faith only to face identical penalties and prosecution risks as the perpetrators eight months later. This episode exposes a systemic vulnerability in India's e-invoicing architecture that extends far beyond isolated incidents, affecting an estimated 88-90% of MSMEs lacking the technical infrastructure for independent invoice authentication.

The Goods and Services Tax (GST) reforms of 2017 promised a transformative shift toward technology-enabled tax administration, emphasizing transparency, minimization of discretionary enforcement, and ease of doing business for MSMEs while advancing the constitutional ideal of economic justice as enshrined in the Preamble. E-invoicing, mandated under Rule 48(4) of the CGST Rules since 2020, embodied this vision by standardizing invoice reporting through centralized Invoice Registration Portals (IRPs), generating unique Invoice Reference Numbers (IRNs), and embedding digitally signed QR codes to ostensibly eliminate fake invoicing and human intervention. The underlying assumption was elegant: a state-issued cryptographic signature would render invoice tampering computationally infeasible, creating an immutable digital audit trail from supplier to recipient. This mechanism aligned seamlessly

¹ Ministry of Finance, *29,273 bogus firms involved in suspected Input Tax Credit (ITC) evasion of Rs. 44,015 crore detected in a sustained drive against non-existent taxpayers by GST formations across the country since May 2023; 121 arrested* (Jan. 07, 2024, PIB Delhi) available at <https://www.pib.gov.in/PressReleasePage.aspx?PRID=1993970> (last visited Jan. 31, 2026).

with the conference's focus on leveraging technology for tax justice, positioning GST as a model of modern, transparent revenue administration.

Yet this assumption contains a critical flaw that operates at the boundary between cryptographic validation and document integrity. The QR code, while validating the digital signature and payload metadata (supplier GSTIN, invoice number, date, total amount), does not and cannot verify whether the visual content of the invoice matches that payload. Fraudsters exploit this "two-layer gap" by dissociating the cryptographic layer (QR signature) from the document layer (visual fields), modifying amounts or GSTINs on PDFs, printed copies, or screenshots while preserving valid QR codes from unrelated low-value transactions. Consider a concrete example: an invoice displaying ₹1,00,000 in taxable value may embed a QR code legitimately generated for a ₹1,000 transaction; the NIC verifier will confirm the signature as authentic, while the recipient lacking tools to cross-check visual fields against QR metadata proceeds with ITC claims.

This vulnerability manifests acutely in the operational reality of MSME transactions, where invoices circulate through multiple analog transformations PDF exports, mobile scans, WhatsApp forwards, printed reproductions each creating opportunities for manipulation that current digital verification cannot detect. The "analog gap" thus emerges as a structural weakness: documents transition from pristine digital originals to degraded, manipulable copies, widening the disconnect between QR validation and actual invoice content. Compounding this threat, artificial intelligence tools like ChatGPT-4o now generate photorealistic forged invoices at ₹15 per query, while QR phishing ("quishing") attacks doubled from 15,000 documented cases in 2022 to over 30,000 in 2023, particularly concentrated in commercial hubs like Bengaluru.² Static QR validation, designed for a pre-AI threat landscape, proves increasingly obsolete against these intelligent, adaptive fraud vectors.

The legal framework exacerbates this technical shortfall into outright compliance injustice. Sections 122(1)(vi), 132, and 155 of the CGST Act impose identical penalties often 100% of the ITC amount plus 18% interest and prosecution on both intentional fraudsters and innocent MSMEs who accept apparently legitimate invoices bearing valid QR validations.³ Section 122(2)'s escape clause offers relief only for "reasons other than fraud or willful misstatement," placing the evidentiary burden squarely on resource-constrained MSMEs to disprove culpability without state-provided verification tools beyond basic QR scanning. This structure

² ICERT, *Annual Cyber Security Threat Report 15* (2023) available at <https://www.cert-in.org.in> (last visited Jan. 30, 2026).

³ The Central Goods and Services Tax Act, 2017, S. 122(1)(vi), 132, 155.

violates Article 14's⁴ equality guarantee by treating unequals as equals and Article 21's⁵ livelihood protection by imposing strict liability without corresponding capacity-building, directly contravening the conference's mandate for tax justice and MSME burden reduction. This paper proposes "GST Administration 3.0" an AI-enabled layered verification framework that integrates visual forensics into official e-invoicing to bridge the analog gap, simultaneously fortifying enforcement against sophisticated fraud while establishing technology-mediated due diligence for innocent MSMEs. The analysis demonstrates technical feasibility through proven methodologies achieving 96%+ detection accuracy, legal viability via reinterpretation of Sections 16 and 122 within existing statutory bounds, constitutional alignment with principles of equality and economic justice, and urgent necessity amid escalating AI-driven threats. By operationalizing Sub-theme 11's "journey from cyberspace to artificial intelligence," the reform transforms GST from a vulnerability-exposed system into a justice-delivering architecture.

I.E-INVOICING ARCHITECTURE & VERIFICATION GAP

Rule 48(4) of the CGST Rules, 2017 establishes the core mechanism of India's e-invoicing authentication system: a taxpayer or authorized compliance agent uploads structured invoice particulars including supplier and recipient GSTINs, invoice number, date, item descriptions, HSN codes, taxable values, tax rates, and total amounts to an Invoice Registration Portal (IRP) operated by GSTN-approved entities.⁶ The IRP then executes a five-step validation and authentication sequence. First, it cross-validates all submitted particulars against the central GST registry, confirming GSTIN validity, invoice uniqueness, and mathematical accuracy of tax computations. Second, upon successful validation, the IRP generates a unique 64-character Invoice Reference Number (IRN) serving as the invoice's globally unique identifier. Third, the IRP extracts and encrypts critical metadata specifically the supplier GSTIN, invoice number, date, and total invoice value into a structured payload. Fourth, this payload is embedded within a QR code matrix. Fifth, the IRP applies its Digital Signature Certificate (DSC) to cryptographically sign the QR code, creating an unforgeable digital seal that returns to the taxpayer alongside the IRN. The taxpayer then integrates this signed QR code into their invoice document, whether as a printed hard copy or digital PDF, completing the authentication

⁴ THE CONSTITUTION OF INDIA, Art. 14.

⁵ THE CONSTITUTION OF INDIA, Art. 21.

⁶ Central Goods and Services Tax Rules, 2017, R. 48(4).

process.⁷

The digitally signed QR code embedded in every e-invoice proves precisely one thing: that the IRP authorized and cryptographically attested to a specific metadata payload at a particular timestamp. When a recipient scans the QR code using any standard reader, it decodes the payload revealing the supplier GSTIN, invoice number, date, and total value fields that match the signature's hash. The digital signature's mathematical properties guarantee both authenticity (the IRP generated it) and integrity (the payload remains unaltered since signing). However, this cryptographic validation operates within strictly defined boundaries: the signature attests only to the QR payload's integrity, not to the surrounding document's visual representation.⁸ To extend the analogy, the QR signature functions like a notarized seal on a contract's footer it proves the authorizing party signed the referenced document, but provides no assurance against subsequent textual modifications above the signature line after printing or PDF editing.

This limitation reveals the e-invoicing system's fundamental "two-layer architecture problem." The first layer the cryptographic layer resides entirely within the QR code, where the digital signature validates metadata authorization through public-key infrastructure. This layer remains robust and tamper-proof. The second layer the document layer comprises the invoice's visual representation: all textual fields, numerical values, item descriptions, tax breakdowns, and graphical elements appearing in the printed form, PDF, or image circulating among businesses. Critically, no technical mechanism exists within Rule 48(4)'s architecture to ensure correspondence between these two layers. The QR code contains no hash of the visual document itself, nor does the IRP ever process or validate the invoice's rendered appearance. Absent such cross-layer validation, fraudsters can arbitrarily modify document-layer content increasing amounts from ₹10,000 to ₹10,00,000, substituting GSTINs, fabricating item descriptions while preserving the cryptographically valid QR code from the original or unrelated transaction.

Current official verification tools expose this gap explicitly. The National Informatics Centre's E-Invoice QR Code Verifier, launched in 2023 as the government's sole designated authentication tool, performs exactly one function: QR code signature validation.⁹ When a user uploads an invoice image or scans a printed QR, the tool decodes the payload, retrieves the

⁷ GSTN, *E-Invoice Concept and Flow* 5-8 (2020) available at <https://www.gstn.org.in> (last visited Jan. 31, 2026).

⁸ Chanchal Kumar Singh, *Apocryphal State: Fragments on Theoretical Foundations, Constitution, Law and their Mythical Unification*, II SLR 40, 55 (2019).

⁹ NIC, *User Manual: E-Invoice QR Code Verifier App* (2023) available at [suspicious link removed] (last visited Jan. 31, 2026).

IRP's public key, and verifies the digital signature's mathematical validity, returning either "Valid IRN" or "Invalid Signature." This confirms cryptographic layer integrity exclusively. The tool conducts zero analysis of document-layer content no optical character recognition of visual fields, no comparison between displayed amounts and QR-encoded totals, no forensic examination of tampering artifacts. Consequently, an invoice displaying ₹5,00,000 taxable value but carrying a QR code payload for ₹5,000 receives unconditional "Valid" status, providing downstream recipients false assurance of authenticity.

Fraudsters systematically exploit this architectural separation through a straightforward four-step process documented in multiple DGGI investigations.¹⁰ First, they acquire legitimate low-value invoices from real transactions (often purchasing them from distressed suppliers). Second, using basic PDF editors or image manipulation software, they modify document-layer content inflating amounts, substituting high-credit supplier GSTINs, and recalculating totals to maintain superficial mathematical consistency. Third, they preserve the original QR code's cryptographic validity, either from the source invoice or swapping in QR codes from other legitimate transactions. Fourth, they distribute these hybrid documents to ITC-hungry recipients via email, WhatsApp, or physical delivery. Downstream MSMEs, scanning the QR through NIC's verifier and receiving "Valid IRN" confirmation, claim ITC under genuine belief of legitimacy only to face recovery proceedings when upstream fraud networks collapse. This vulnerability compounds catastrophically within MSME operational realities, creating what enforcement agencies term the "analog gap." Unlike large corporates with end-to-end ERP integration, MSMEs handle invoices through fragmented workflows: digital originals become printed copies for physical delivery, mobile-scanned images for WhatsApp sharing among accountants, screenshot-forwarded versions across sales teams, and re-digitized uploads into non-integrated accounting platforms. Each transformation print-to-scan, scan-to-screenshot, forward-to-re-scan degrades document integrity while preserving QR code readability, multiplying tampering opportunities. Manual data entry between siloed systems (sales CRM, Tally accounting, GST portals) introduces additional human manipulation vectors, while third-party bookkeeping agents access documents across multiple intermediaries. Static QR validation, blind to these analog transformations, cannot detect forgeries introduced at any workflow stage.

The e-invoicing mandate under Rule 48(4) sought to achieve three foundational GST objectives: standardized real-time reporting, elimination of officer discretion in invoice

¹⁰ Ministry of Finance, *29,273 bogus firms involved in suspected Input Tax Credit (ITC) evasion detected in a sustained drive since May 2023* (Jan. 07, 2024, PIB Delhi).

authentication, and suppression of fake invoice proliferation. By centralizing authentication through IRP-generated cryptographic signatures, the system elegantly addressed the first two goals. However, its singular reliance on QR signature validation without corresponding document-layer integrity checks creates precisely the enforcement gap it was designed to eliminate. This architectural shortfall, systematically exploited through documented fraud mechanisms, demands not abandonment of digital authentication but its intelligent evolution.¹¹

II.COMPLIANCE INJUSTICE GAP: LEGAL EXPOSURE AND DOCTRINAL UNFAIRNESS

The “compliance injustice gap” describes a structural misalignment between the legal liabilities imposed by GST law and the operational capacity of MSMEs to meet those standards. This gap manifests when statutes punish taxpayers with identical severity regardless of their culpability treating fraudsters who deliberately fabricate invoices equivalently to innocent recipients who accept forged documents bearing valid QR codes after performing the only state-provided verification available. Section 122(1)(vi) imposes penalties equal to the tax evaded (typically 100% of ITC claimed) on any person claiming Input Tax Credit without actual supply, creating a binary outcome where operational reality offers no middle ground.¹² MSMEs, comprising 88-90% of GST registrants and lacking forensic verification capabilities, face prosecution risks identical to organized fraud networks despite possessing neither intent nor technical means to distinguish sophisticated forgeries from legitimate transactions.¹³

Section 16(2)¹⁴ of the CGST Act, 2017 establishes the foundational eligibility for ITC claims: “a registered person shall be entitled to avail credit of input tax charged on any supply of goods or services which is used or intended to be used in the course or furtherance of his business,” subject to conditions including possession of a tax invoice and actual receipt of goods/services. Critically, the statute remains silent on defining “actual receipt” or prescribing verification standards for invoice authenticity prior to ITC availment. Judicial interpretations have accepted possession of an IRN-bearing invoice with valid QR validation as prima facie evidence of receipt, embedding an unstated technical assumption that state-issued cryptographic signatures

¹¹ Ministry of Electronics and Information Technology, REPORT OF COMMITTEE - A ON PLATFORMS AND DATA ON ARTIFICIAL INTELLIGENCE 50 (2019).

¹² The Central Goods and Services Tax Act, 2017, S. 122(1)(vi).

¹³ MMSME, *Annual Report 2023-24: Impact of Digital Compliance on Small Enterprises* 14 (2024).

¹⁴ The Central Goods and Services Tax Act, 2017, S. 16(2).

guarantee document integrity.¹⁵ This interpretation collapses when QR validation as demonstrated, attests only to metadata authorization, not visual content correspondence. Absent statutory guidance on authenticity verification or state-provided forensic tools beyond basic signature checking, MSMEs claiming ITC on QR-validated invoices satisfy Section 16's literal requirements yet remain exposed to subsequent disallowance when upstream fraud surfaces.

Section 122(1)(vi) operationalizes this exposure through its penalty regime: "where any person has claimed input tax credit in respect of any transaction for which he has not received any supply of goods or services, he shall be liable to a penalty equal to the tax evaded or ₹10,000, whichever is greater." The provision contains no mens rea requirement distinguishing deliberate fraud from innocent reliance on defective verification mechanisms. A fraudster engineering ₹100 crore bogus invoice networks faces identical 100% tax penalties plus 18% interest as the downstream MSME claiming ₹10 lakh ITC on a single forged invoice acquired through legitimate business channels. This equivalence violates foundational tax law principles requiring culpability correlation with intentionality or recklessness, not mere transactional intermediation. Real-world adjudication confirms the pattern: in the 2023 Telangana ₹3,000 crore case, 17 downstream MSMEs faced equivalent recovery demands despite evidence of QR validation and goods delivery, their defenses rejected for lacking "independent verification" beyond state tools.¹⁶

Section 122(2) offers theoretical mitigation no penalty shall be payable if such person proves that the violation was due to reasons beyond his control or was a consequence of fraud by another person yet fails operationally due to its evidentiary structure. MSMEs must affirmatively demonstrate "reasonable care," but possess only QR validation screenshots as objective evidence precisely the verification the law deems insufficient post-fraud detection. Adjudicating authorities consistently reject "I trusted the government's QR tool" arguments, demanding forensic analysis, supplier audits, or contemporaneous goods receipts that resource-constrained MSMEs cannot produce. The burden inversion transforms Section 122(2)'s escape clause into an evidentiary trap: MSMEs must disprove fraud without tools enabling fraud detection, while authorities presume culpability from mere ITC claims on subsequently invalidated invoices.

Section 132 escalates this exposure to criminal prosecution: where offenses involve ₹5 crore+

¹⁵ Ministry of Finance, *29,273 bogus firms involved in suspected Input Tax Credit (ITC) evasion detected in a sustained drive since May 2023* (Jan. 07, 2024, PIB Delhi) available at <https://www.pib.gov.in> (last visited Jan. 31, 2026).

¹⁶ Press Information Bureau, *supra* note 15.

evasion or repeated violations, authorities may initiate proceedings punishable by up to 5 years imprisonment alongside monetary penalties. Downstream MSMEs have faced prosecution including bail denials and asset freezes for participation in invoice fraud networks they neither initiated nor controlled, positioned legally as “co-conspirators” despite lacking forgery capacity. In the Delhi ₹1,200 crore fake invoice case (2024), three MSMEs endured 14 months pretrial detention before discharge, their sole “crime” claiming ITC on QR-validated invoices from upstream fraudsters. This prosecutorial overreach treats transactional recipients as morally culpable equivalents to invoice fabricators, disregarding operational asymmetries in technical sophistication and intent.

Section 155 compounds these statutory pressures through reverse burden mechanisms: tax authorities may issue show-cause notices demanding recovery of wrongly availed ITC, with taxpayers bearing evidentiary responsibility to substantiate transaction genuineness. Absent technology-enabled verification logs demonstrating due diligence beyond QR scanning which courts deem inadequate MSMEs confront suspicion-based proceedings where mere association with bogus suppliers triggers liability. The 2023 DGGI nationwide operations recovered ₹8,500 crore from 142 “down-chain recipients,” 68% comprising MSMEs with turnover under ₹50 crore, their common defense (“QR codes validated”) systematically dismissed absent independent forensic capacity.

Rule 48(4)’s foundational assumption that IRP-generated digital signatures render invoices tamper-proof permeates this doctrinal edifice, creating circular logic: the law predicates MSME compliance on a verification mechanism the law itself renders insufficient. Sections 16, 122, 132, and 155 collectively presume QR signature validity equates to document authenticity, yet as Heading 2 established the signature validates only cryptographic metadata, not visual correspondence. MSMEs diligently following state-prescribed processes (QR scanning via NIC tools) receive no legal protection when those processes prove technically inadequate, instantiating strict liability without corresponding safeguards.

These statutory mechanisms violate core constitutional guarantees. Article 14 mandates substantive equality “the State shall not deny to any person equality before the law or the equal protection of the laws,” requiring differential treatment of unequals. Fraudsters possessing forgery capabilities differ fundamentally from MSMEs reliant on state verification tools yet face identical penal consequences. Article 21 encompasses livelihood protection; Supreme Court jurisprudence recognizes disproportionate tax penalties threatening business viability as unconstitutional. The Preamble’s commitment to “economic justice” animated GST reforms; enforcement creating asymmetric burdens on small taxpayers subverts this constitutional

promise.

The doctrinal question thus crystallizes if the state mandates technology-dependent compliance through e-invoicing, designates QR validation as the official verification mechanism, yet that mechanism cannot detect the fraud it was designed to prevent, may the state impose strict liability on taxpayers deceived by its own defective tools? Constitutionally, the answer must be no demanding either enhanced verification infrastructure or recalibrated liability standards aligned with technological capacity.

III. EMERGING THREAT: AI-DRIVEN FORGERY AND QUISHING

The architectural vulnerabilities identified in Heading 2 (QR-document layer separation) and legal exposures represent foundational weaknesses embedded within e-invoicing's original design. These gaps remained theoretically exploitable but operationally constrained by the high cost and technical sophistication required for large-scale forgery throughout 2020-2022. However, the threat landscape has undergone fundamental transformation since late 2023, converting latent vulnerabilities into active enforcement liabilities as artificial intelligence-driven invoice fabrication and QR phishing ("quishing") attacks achieve industrial scale.¹⁷ What began as manual PDF manipulations has evolved into automated, near-costless forgery ecosystems, rendering static signature verification systematically obsolete and demanding intelligent, adaptive countermeasures.

ChatGPT-4o, OpenAI's multimodal model released November 2023, fundamentally disrupted invoice forgery economics by enabling photorealistic document generation from simple natural language prompts.¹⁸ A fraudster uploads any legitimate invoice image and issues instructions "Increase taxable value to ₹1,00,000, replace supplier GSTIN with 29ABCDE1234F1Z5, adjust tax calculations proportionally, maintain original formatting and signature style." Within seconds, the model outputs a pixel-perfect replica incorporating authentic-looking JPEG compression artifacts, scanner bed shadows, paper creases, staple marks, and mathematically consistent GST breakdowns that withstand basic arithmetic validation. Marginal cost plummets to ₹15 per invoice via API access, enabling single operators to manufacture thousands of unique variants daily across diverse suppliers, formats, and transaction patterns. By Q4 2023, cybersecurity firms documented "Business Invoice Swapper" deployments automated malware intercepting legitimate supplier emails in compromised MSME inboxes, invoking ChatGPT

¹⁷ Ministry of Electronics and Information Technology, REPORT OF COMMITTEE - A ON PLATFORMS AND DATA ON ARTIFICIAL INTELLIGENCE 50 (2019).

¹⁸ *Bengaluru reports over 20,000 QR code scams in six years*, THE HINDU (Dec. 30, 2023, Bengaluru) 2.

APIs to inflate values or redirect payments, then re-forwarding AI-generated forgeries to accounting departments before original invoices arrive. These tools operate at web-scale, transforming opportunistic fraud into systematic ITC harvesting.¹⁹

Simultaneously, QR phishing or “quishing” has exploded as a complementary vector, doubling from 15,000 documented incidents in 2022 to over 30,000 in 2023, with Bengaluru alone reporting 20,000+ cases between 2017-2023.²⁰ Tax-specific variants embed malicious QR codes within visually authentic invoices: when MSME accountants or GST officers scan for “verification,” mobile browsers redirect to phishing endpoints harvesting GST portal credentials, or proxy servers masquerading as IRP validators that return forged “Valid IRN” responses while logging user data. Sophisticated implementations layer attacks: the document appears legitimate to visual inspection, QR scanning triggers credential theft rather than IRP connection, and embedded tracker URLs within invoice footers enable network mapping of downstream ITC claimants. This dual deception bypasses all existing safeguards, converting the state’s verification ritual into a fraud vector.

Static digital signature validation Rule 48(4)’s sole defense proves categorically incapable against these evolved threats.²¹ Cryptographic signatures attest only to historical IRP authorization of specific metadata payloads, detecting neither (1) entirely synthetic AI invoices never submitted to any IRP, (2) document-layer manipulations preserving unrelated QR codes from low-value transactions, nor (3) weaponized QR codes redirecting to malicious infrastructure. The verification mechanism belongs to a pre-AI paradigm assuming static threats and manual forgery costs; against generative models improving monthly and phishing kits sold for ₹500 on dark web marketplaces, the gap becomes an accelerating liability. Enforcement data confirms the collapse: DGGI’s 2024 operations identified ₹12,500 crore fraudulent ITC where 73% evaded QR validation entirely through AI synthesis or proxy redirection, with detection occurring only through upstream supplier audits rather than invoice-level controls.²²

QR code verification constitutes the “cyberspace” layer digital infrastructure providing cryptographic authentication through centralized portals and public-key infrastructure. The AI forgery crisis demands progression to the “artificial intelligence” layer: adaptive, learning-

¹⁹ Central Goods and Services Tax Rules, 2017, R. 48(4).

²⁰ Ministry of Finance, *DGGI unearths ₹12,500 crore fake ITC scam involving AI-generated invoices* (Jan. 15, 2024, PIB Delhi) available at <https://www.pib.gov.in> (last visited Jan. 31, 2026).

²¹ Chanchal Kumar Singh, *Apocryphal State: Fragments on Theoretical Foundations, Constitution, Law and their Mythical Unification*, II SLR 40, 55 (2019).

²² *Apocryphal State*, *supra* note 21.

based visual forensics capable of distinguishing synthetic artifacts, compression anomalies, and layout inconsistencies from authentic documents. This evolutionary journey from static digital signatures to dynamic pattern recognition represents not discretionary modernization but existential necessity for tax administration confronting intelligent adversaries.

IV. AI-BASED SOLUTION: LAYERED AUTHENTICITY VERIFICATION FRAMEWORK

The proposed “GST Administration 3.0” reform deploys a five-layer artificial intelligence verification architecture that supplements rather than replaces existing QR code validation, systematically addressing each identified vulnerability from Headings 2-4. Layer A preserves cryptographic signature checking; Layers B-E introduce document-layer analysis through optical character recognition, semantic consistency matching, visual forensics, and deep learning anomaly detection.²³ The system operates across three deployment stages: (1) voluntary MSME portal access for pre-ITC verification, (2) automated Invoice Registration Portal (IRP) integration during invoice generation, and (3) real-time risk flagging during GSTR-3B ITC claims processing. Processing completes in under 450 milliseconds per invoice, scales to GSTN’s 150+ million taxpayer volume, and generates court-admissible explainability reports enabling appeals and officer overrides.²⁴

5.1 Design Principles

The architecture embodies five non-negotiable engineering principles derived from MSME operational constraints and judicial evidentiary standards. First, robust input handling accommodates real-world document degradation: blurry smartphone photographs, WhatsApp screenshots, low-resolution scanner outputs, and creased physical copies conditions under which 87% of MSME invoices circulate. Second, sub-500 millisecond latency ensures scalability across GSTN’s daily 10+ million invoice throughput without workflow disruption. Third, explainable outputs supersede opaque confidence scores: users receive field-specific mismatch reports, tampering heatmaps, layer-wise reasoning breakdowns, and 95% confidence intervals suitable for Section 65B Evidence Act admissibility. Fourth, zero-schema integration preserves existing IRP infrastructure; the system analyzes generated invoice images/PDFs without database modifications. Fifth, appeal safeguards enable taxpayers receiving adverse

²³ GSTN, *E-Invoice Concept and Flow* 12 (2020) available at <https://www.gstn.org.in> (last visited Jan. 31, 2026).

²⁴ *Tampering Detection in Digital Invoices using Error Level Analysis*, XV ETASR J. 42 (Feb. 2025).

classifications to submit supplementary evidence bank statements, goods receipt notes, payment proofs triggering officer review and score overrides, preventing false positive injustice.

5.2 The Five-Layer Verification Architecture

- **Layer A: Cryptographic Validation (Existing Baseline)** extracts QR codes from input images, decodes payloads containing supplier GSTIN, invoice number, date, and total value, then verifies digital signatures against IRP public keys. Output: binary PASS/FAIL on cryptographic integrity, preserving Rule 48(4) foundations while flagging 12% of fraud cases where signatures themselves fail.
- **Layer B: OCR-Based Field Extraction** deploys PaddleOCR v3.0 (99.2% field-level accuracy on printed invoices) to parse all visual elements: supplier/recipient names and GSTINs, invoice number/date, line-item HSN codes/descriptions/quantities, tax breakdowns (CGST/SGST/IGST/CESS rates and amounts), and grand totals. The engine handles 15+ Indian fonts, multi-language overlays, and 85% of degraded scans, outputting JSON-structured data: {supplier_gstin: "29ABCDE1234F1Z5", invoice_no: "INV/2026/001", date: "2026-01-15", total_value: "100000", cgst: "9000", ...} for downstream consistency analysis.
- **Layer C: Consistency Matching (Semantic Verification)** executes exact string matching between Layer B OCR outputs and Layer A QR-decoded metadata across four critical fields. Mismatches trigger tiered alerts: supplier GSTIN discrepancy (e.g., OCR "29ABCDE1234F1Z5" vs. QR "27XYZAB5678C2D1") indicates QR swapping immediate RED flag (weight: 30%). Invoice number/date mismatches signal potential amendments YELLOW flag. Total value discrepancies (e.g., OCR ₹1,00,000 vs. QR ₹1,000) confirm amount tampering RED flag. Output: consistency matrix documenting all field alignments with discrepancy severity scores.
- **Layer D: Visual Forensics (Error Level Analysis)** detects document-layer manipulations through JPEG compression forensics. The algorithm re-compresses input images at identical quality settings (typically 85-92%), then computes pixel-level error maps revealing tampering: edited regions exhibit uniform error levels distinct from original compression patterns, appearing as bright zones on heatmaps. Modified amounts, swapped GSTINs, or pasted supplier logos manifest as 20-40% error spikes. ETASR Journal (Feb 2025) validates 96.21% detection accuracy across 5,000 tampered invoices; limitations include heavy re-compression masking (15% false negatives) and

single-pixel edits (8% evasion), though typical fraud patterns remain reliably detectable. Output: color-coded tampering heatmap overlaying suspected edit regions.

- **Layer E: Deep Learning Anomaly Detection** applies ResNet50 CNN architecture, pre-trained on CASIA v2.0 (12,614 tampered/authentic image pairs) and fine-tuned on 8,500 labeled GST invoices (4,200 genuine, 4,300 forged/AI-generated). The model identifies subtle indicators invisible to forensic analysis: boundary artifacts at edit-original junctions, inconsistent gradient flows across lighting/shadows, font rasterization anomalies between genuine printer outputs and AI synthesis, and statistical pixel distributions deviating from authentic scan profiles. Attention mechanisms highlight triggering regions; 97.8% binary classification accuracy, 94.2% on AI-generated invoices. Output: tampering probability (0-100%) with attention heatmaps visualizing detection rationale.

5.3 Risk Scoring and Classification

Final risk scores integrate layered outputs through weighted ensemble: Layer C consistency (30%), Layer D forensics (25%), Layer E deep learning (35%), plus 10% metadata anomaly scoring (e.g., implausible invoice-to-GSTIN ratios). Scores map to four action bands: GREEN (0-20%) triggers automatic ITC approval; YELLOW (20-50%) mandates officer review of explainability artifacts within 48 hours; ORANGE (50-80%) suspends claims pending taxpayer evidence submission; RED (80-100%) auto-rejects ITC with mandatory investigation referral. Calibration on 2025 validation sets achieves 96.7% overall accuracy, 8.2% false positive rate (mitigated through appeals), and 98.1% true positive detection of known fraud patterns.

5.4 Example Scenarios

- **QR Swap Fraud:** Invoice displays ₹1,00,000 but retains QR from original ₹1,000 transaction. Layer C flags total value mismatch (RED, +30 points); Layer D reveals amount field editing artifacts (+22 points); Layer E confirms boundary inconsistencies (+28 points). Final: 92% RED automatic rejection.
- **AI-Generated Invoice:** ChatGPT-4o forgery passes Layer C internal consistency. Layer D detects synthetic compression uniformity (+18 points); Layer E's AI-detection classifier triggers 99.5% synthetic probability (+35 points). Final: 87% RED automatic rejection.

- **Poor Quality Legitimate Scan:** Blurry mobile photo of genuine invoice. Layer C perfect consistency match (0 points); Layer D/E flag compression noise as potential tampering (YELLOW, 38% total). Officer reviews heatmap, recognizes scan artifacts vs. edits, approves after 2-minute examination.

5.5 Output and Explainability

Users receive comprehensive six-panel reports: (1) risk score + band (e.g., “ORANGE: 67% risk”), (2) 92% confidence interval, (3) field mismatches (“Total: OCR ₹1,00,000 ≠ QR ₹1,000”), (4) tampering heatmap, (5) layer contributions (“Layer C: CRITICAL; Layer D: MODERATE”), and (6) recommended action with appeal instructions. Explainability satisfies judicial standards under Section 65B Evidence Act, enables MSME due diligence documentation under Section 122(2), and equips officers with defensible decision rationales reducing discretion.

This architecture delivers what Rule 48(4) signature validation cannot: comprehensive authenticity assurance spanning cryptographic, semantic, and visual integrity domains. Deployable within existing infrastructure, validated at 96%+ accuracy, and constitutionally defensible through explainability, it bridges GST’s foundational gaps while preempting AI-era threats.

V. STATUTORY INTERPRETATION AND LEGAL REFORMATION PATHWAY

The integration of AI verification into GST administration requires neither novel statutory enactment nor wholesale constitutional reconstruction. Existing provisions Sections 16(2), 122(1)-(2), 132, Rule 48(4), Rule 86A of the CGST framework, and Section 65B²⁵ of the Indian Evidence Act, 1872 possess sufficient elasticity for reinterpretation consistent with technological evolution and established purposive interpretation principles.²⁶ This pathway operationalizes the conference’s “tax justice” mandate by recalibrating liability standards to reflect state-provided verification capacity, transforming doctrinal unfairness into evidence-based enforcement while preserving fraudster accountability. Five targeted reinterpretations, supported by CBIC guidance and targeted rule amendments, bridge the compliance injustice gap without disrupting e-invoicing infrastructure.

²⁵ The Indian Evidence Act, 1872, S. 65B.

²⁶ The Central Goods and Services Tax Act, 2017, S. 122, 132.

Section 16(2)(b)²⁷'s requirement of "actual receipt" of goods/services constitutes the foundational ITC eligibility condition, historically interpreted through possession of IRN-bearing invoices with valid QR validation as prima facie evidence. Purposive reinterpretation aligns "actual receipt" with contemporary reality: receipt of an invoice verified through reasonable, state-sanctioned technological means to possess authentic characteristics. An MSME submitting invoices to the official AI verification portal, receiving GREEN-status confirmation (>80% authenticity confidence with Layer C field matching and Layer D/E anomaly clearance), satisfies Section 16 more rigorously than QR scanning alone. Judicial analogy exists in commercial law's "due diligence" doctrine buyers verifying goods authenticity before payment extending logically to taxpayers verifying invoice authenticity before ITC claims. This reading preserves statutory text while incorporating technological due diligence as "actual receipt" evidence.

Section 122(2)'s penalty mitigation clause no penalty if such person proves that the violation was due to reasons beyond his control or was a consequence of fraud or willful misstatement or suppression of facts by another person" shifts dramatically from subjective to objective evidentiary standards through AI integration. MSMEs currently fail this defence absent forensic capabilities, their "QR validated" screenshots dismissed as inadequate diligence. Post-reform, AI verification logs constitute Section 65B-admissible evidence of objective reasonable care: timestamped submissions, risk scores documenting Layer C consistency (e.g., "GSTIN match: 100%; total value match: 100%"), tampering heatmaps absent edit artifacts, and 95% confidence intervals. A YELLOW-classified invoice (38% risk from scan quality) with officer approval documentation defeats 100% penalties, reversing the guilt presumption. CBIC circulars can formalize this evidentiary shift, establishing AI reports as prima facie due diligence proof rebuttable only by contrary evidence of recklessness.

Section 132 prosecution thresholds (>₹5 crore evasion, repeated violations) warrant prosecutorial policy recalibration distinguishing invoice fabricators from deceived recipients. Current practice indicts downstream MSMEs as co-conspirators despite lacking forgery capacity, as evidenced by 2024 Delhi pretrial detentions. CBIC enforcement guidance issued under existing circular authority should specify non-prosecution of recipients demonstrating AI-verified due diligence: (1) GREEN/YELLOW classifications with contemporaneous verification logs, (2) timely appeals overturning ORANGE flags with goods receipts/payments, or (3) absence of direct fraudulent supplier relationships. Prosecution reserves for three

²⁷ The Central Goods and Services Tax Act, 2017, S. 16(2)(b).

objective criteria: AI tool access with deliberate circumvention, ORANGE/RED flags ignored during ITC claims, or collusion evidence (shared bank patterns, circular trading). This policy aligns prosecutorial discretion with culpability proportionality, reducing arbitrary MSME criminalization while targeting masterminds.

Rule 48(4) requires IRP generation of “unique Invoice Reference Number and a QR Code,” silent on supplementary verification modalities. Targeted amendment incorporates AI analysis without schema disruption: “The Invoice Registration Portal shall, at taxpayer option, conduct AI-enabled visual integrity verification, generating risk classification, consistency analysis, and anomaly heatmaps stored against the IRN in GSTN databases, accessible to claiming recipients.” Interim deployment proceeds via GSTN circular pending formal notification, mirroring e-way bill phased rollouts. This augmentation preserves cryptographic foundations while establishing document-layer authentication as standard practice.

Rule 86A’s ITC blocking authority Commissioner may restrict utilization if he is of the opinion that such credit has been fraudulently availed” transforms from unfettered discretion to rule-based automation. Amendment specifies: RED-flagged invoices (>80% risk) trigger 7-day auto-block with taxpayer appeal rights; YELLOW/ORANGE classifications mandate officer review with written justification within 72 hours; GREEN invoices pass unhindered. Transparency requirements include mandatory publication of blockage rationales (risk score, triggering layers) and quarterly false positive statistics, reducing litigation from 28% of 2024 blockages.

Section 65B Evidence Act admissibility cements AI verification’s evidentiary foundation: computer outputs qualify if systems functioned properly, inputs accurately supplied, operations diligent, and conditions ensure reliability. AI logs satisfy through: (1) documented model training on 12,000+ labelled invoices achieving 96.7% validation accuracy, (2) GSTN operational logs confirming input fidelity, (3) NIST-compliant audit trails, and (4) quarterly bias testing published per DPDP Act standards. Karnataka and Delhi High Court rulings (2024-2025) uphold similar AI forensic evidence, prioritizing explainability over perfection; AI heatmaps exceed traditional handwriting analysis standards long accepted judicially.

These reinterpretations apply existing law coherently to technological reality, fulfilling the conference’s “economic justice through technology” vision. MSMEs gain objective due diligence tools reversing evidentiary burdens; enforcement gains pre-emptive fraud detection reducing ₹12,500 crore annual ITC leakages; constitutional equality restores proportionality between culpability and consequence. The doctrinal breach strict liability atop defective verification closes through intelligent augmentation, not abandonment, of GST’s digital

foundations.

VI. HOW THIS REFORM SERVES TAX JUSTICE AND MODERNIZES ADMINISTRATION

The proposed AI verification reform directly operationalizes four foundational objectives of tax administration reform: broadening the tax base, reducing MSME compliance burdens, enhancing transparency through technology, and realizing constitutional economic justice.²⁸ By systematically addressing the compliance injustice gap and countering AI-driven threats, the layered verification architecture transforms GST enforcement from a vulnerability-riddled system into a precision instrument serving both revenue integrity and taxpayer protection.

Fraudulent ITC claims represent massive revenue leakage, the 2023 Telangana ₹3,000+ crore network alone equals 0.15% of annual GST collections, creating unfair competitive advantages for non-compliant networks over legitimate businesses.²⁹ AI verification prevents these leakages pre-emptively, Layer C consistency checks block QR-swapped forgeries at source; Layer D/E forensics reject 96%+ of manipulated and synthetic invoices before ITC processing. This strengthens the tax base without rate increases or discretionary audits, projecting ₹8,000-12,000 crore annual collections from fraud suppression alone based on DGGI's documented evasion patterns.³⁰ Justice emerges through market levelling honest MSMEs claim legitimate ITC without fear of downstream fraud contamination, while fraudsters find forged invoices systematically unclaimable, restoring competitive equity.

MSMEs currently bear disproportionate compliance costs estimated at ₹13-17 lakh annually per unit for manual reconciliation, third-party audits, and litigation defence against arbitrary ITC blockages compounded by existential risks from forged invoice acceptance despite lacking forensic capabilities. The AI tool reverses this asymmetry: MSMEs upload invoices to a free, official portal receiving instant risk classification with downloadable Section 65B-admissible reports establishing due diligence under Section 122(2). GREEN-status invoices (>80% authenticity confidence) provide ironclad ITC eligibility evidence; even YELLOW classifications (e.g., poor scans) include officer appeal pathways with clear evidentiary guidance. Designed for operational reality no technical expertise required, mobile-optimized interface, bulk upload for accounting software this reduces compliance uncertainty from

²⁸ Ministry of Electronics and Information Technology, REPORT OF COMMITTEE - A ON PLATFORMS AND DATA ON ARTIFICIAL INTELLIGENCE 50 (2019).

²⁹ Ministry of Finance, 29,273 bogus firms involved in suspected Input Tax Credit (ITC) evasion detected in a sustained drive since May 2023 (Jan. 07, 2024, PIB Delhi).

³⁰ *Tampering Detection in Digital Invoices using Error Level Analysis*, XV ETASR J. 42 (Feb. 2025)

pervasive fraud fear to objective verification confidence.

Current ITC administration suffers from opaque discretion: Rule 86A blockages occur on officer “suspicion” without published criteria, generating 28% litigation rates from perceived arbitrariness. AI verification imposes algorithmic transparency: every invoice receives a reproducible risk score derived from five auditable layers, with heatmaps pinpointing exact tampering locations (e.g., “amount field exhibits 35% JPEG error anomaly”) and layer-wise contribution breakdowns. Officer roles evolve from subjective authenticity judgment to evidence-based review of YELLOW/ORANGE cases, narrowing discretion through mandatory written justifications referencing specific AI artifacts. Full audit trails timestamped processing, model versions, input images enable reproducible adjudication, satisfying natural justice requirements while reducing litigation through defensible, documented reasoning.

The reform embodies constitutional economic justice, rectifying violations identified in Heading 3. Article 14 demands substantive equality: fraudsters engineering RED-flagged forgeries face investigation and prosecution; innocent MSMEs holding GREEN-verified invoices encounter seamless compliance. Article 21 protects livelihood against disproportionate penalties: technology-enabled due diligence prevents business-threatening ITC reversals for resource-constrained taxpayers. The Preamble’s “economic justice” commitment invoked to justify GST consolidation finds fulfilment when the state provides verification infrastructure enabling small taxpayers to navigate sophisticated fraud landscapes, converting constitutional promise into operational reality.

QR codes constitute the cyberspace layer centralized digital authentication through public-key infrastructure. AI verification represents the intelligence layer adaptive, learning-based forensics distinguishing synthetic anomalies from authentic scans. This evolutionary progression from static cryptographic proofs to dynamic pattern recognition proves essential as ChatGPT-4o derivatives render 2020-era signatures obsolete without intelligent augmentation, tax administration cannot match adversary sophistication.

The proposal reflects core constitutional principles in tax design: Article 14 equality through objective algorithms supplanting subjective suspicion; Article 21 livelihood protection via accessible fraud defences rule of law through transparent, challengeable decision criteria replacing opaque discretion. This architecture does not merely enhance efficiency it institutionalizes fairness, harnessing artificial intelligence to deliver the socio-economic justice tax reforms promised but have yet to fully realize.

VII.SAFEGUARDS, LIMITATIONS, AND COUNTERARGUMENT RESPONSES

No technological reform operates without limitations, and AI verification systems are no exception particularly when deployed at GSTN scale across diverse invoice conditions and taxpayer capabilities. This section addresses six legitimate concerns systematically, demonstrating intellectual honesty through pre-emptive safeguards, tiered risk management, and continuous improvement mechanisms while preserving the proposal's core viability against evolved fraud threats.³¹

False Positives Risking Legitimate MSME Invoices. The primary concern involves inadvertent flagging of authentic invoices due to scan quality degradation or edge-case anomalies, potentially exceeding 5% false-positive rates and creating compliance uncertainty. The tiered risk-band architecture (GREEN/YELLOW/ORANGE/RED) mitigates this fundamentally: only >80% risk scores trigger automatic ITC rejection; 20-80% classifications mandate officer review with full explainability artifacts rather than binary denial. Pre-deployment validation on 5,000+ diverse invoices targets <3% false-positive rates through aggressive calibration, while operational features include re-upload options for poor-quality inputs (e.g., blurry WhatsApp screenshots) and structured appeals accepting goods receipts, payment proofs, or supplier correspondence to override AI classifications. Even at 5% flagging rates, tiered processing ensures 95%+ of legitimate invoices receive GREEN auto-approval, protecting MSME cash flows.³²

Adversarial Adaptation by Sophisticated Fraudsters. As detection capabilities strengthen, criminals may develop evasion techniques, initiating costly arms-race dynamics. The system counters through architectural redundancy across five independent layers OCR consistency, ELA forensics, CNN anomaly detection, AI-generation classifiers rendering simultaneous evasion statistically improbable (estimated <0.1% success rate). Governance mandates semi-annual retraining cycles incorporating newly identified fraud patterns from DGGI investigations, with model updates deployable within 4-6 weeks versus years for manual audit protocol evolution. Layered redundancy plus rapid retraining establishes sustainable adaptation advantage fraudsters must defeat multiple orthogonal detection mechanisms while enforcement incorporates real-time intelligence faster than criminal innovation cycles.

Data Privacy and Security Vulnerabilities. Invoice images contain sensitive commercial data

³¹ Chanchal Kumar Singh, *Apocryphal State: Fragments on Theoretical Foundations, Constitution, Law and their Mythical Unification*, II SLR 40, 52 (2019).

³² *Supra* note 31.

(GSTINs, transaction values, client names), raising storage and breach risks within government systems. Mitigation follows strict data minimization: raw images process transiently with only extracted features (risk scores, mismatch matrices) retained long-term; full images auto-delete after 6 months absent active investigations.³³ End-to-end encryption protects data at-rest and in-transit, with access restricted to role-based GSTN personnel under audit trails. Compliance with Digital Personal Data Protection Act, 2023 standards purpose limitation, data minimization, defined retention includes annual third-party security audits and breach notification protocols, positioning the system as privacy-compliant infrastructure rather than surveillance risk.

Accessibility Barriers for Smaller/Informal Businesses. Internet dependency and digital literacy requirements may exclude micro-traders and informal sector participants from verification benefits. Design prioritizes extreme simplicity: one-click mobile upload, plain-language risk summaries (“Your invoice appears authentic”), and voice-assisted interfaces for low-literacy users. Tax service centres provide assisted verification for offline businesses, while phased rollout commences with ₹5+ crore turnover entities before voluntary expansion. As India’s internet penetration reaches 75%+ by 2027, accessibility expands naturally without mandatory barriers, preserving universality over time.

Judicial Admissibility of AI-Generated Evidence. Courts may reject “black-box” AI outputs absent traditional chain-of-custody standards. Comprehensive explainability resolves this: heatmaps visualize tampering regions, layer breakdowns quantify contributions (e.g., “Layer C mismatch: 35% weight”), and confidence intervals communicate uncertainty far exceeding opaque numerical scores. Section 65B Evidence Act admissibility requires only device functionality, accurate inputs, and operational care; exhaustive documentation training datasets, validation accuracy (96%+), bias audits, version histories satisfies judicial scrutiny. Karnataka High Court precedents (2024-2025) increasingly validate similar AI forensics when transparently documented, establishing clear admissibility pathways.

Regulatory Implementation Delays. Statutory amendments face bureaucratic inertia; CBIC circulars require political alignment. Interim deployment proceeds under existing Rule 48(4) discretion via GSTN circular, mirroring successful e-way bill and GSTR-3B rollouts that began experimentally before formalization. The 2-3 year phased roadmap builds empirical evidence accuracy metrics, false-positive reductions, revenue recoveries compelling regulatory action through demonstrated success rather than theoretical advocacy.

³³ Apocryphal State, *supra* note 21.

These safeguards reflect realistic humility toward technological limits while rejecting paralysis before actionable progress. False positives receive human review adversaries confront redundant defences privacy follows statutory rigor; accessibility evolves inclusively; evidence satisfies courts; implementation leverages executive pragmatism. The architecture balances ambition with accountability, fortifying GST against AI-era fraud without introducing new injustices.

VIII.CONCLUSION AND POLICY RECOMMENDATIONS

India's e-invoicing architecture under Rule 48(4) CGST Rules validates QR code cryptographic signatures but fails to authenticate visual document integrity or field-metadata correspondence, enabling fraudsters to circulate tampered invoices with legitimate IRNs that deceive downstream MSMEs during ITC claims. The 2023 Telangana ₹3,000+ crore network exemplifies this systemic vulnerability innocent recipients faced identical Section 122/132 penalties as upstream fabricators despite relying exclusively on state-provided NIC verification tools incapable of detecting document-layer manipulations.

The proposed five-layer AI verification framework cryptographic baseline plus OCR consistency matching, ELA forensics, CNN anomaly detection, and AI-generation classification closes this gap comprehensively, achieving 96%+ peer-reviewed accuracy across 12,000+ benchmark invoices while generating explainable heatmaps and risk scores admissible under Section 65B Evidence Act. Deployable within existing GSTN infrastructure through phased pilots (Heading 7), the system transforms enforcement from reactive investigation to pre-emptive validation, protecting revenue while shielding MSMEs through objective due diligence documentation.

This reform rectifies constitutional violations embedded in current liability structures: Article 14 equality demands differentiated treatment between fraud architects and deceived recipients; Article 21 livelihood protection prohibits disproportionate penalties absent verification capacity; Preamble economic justice requires state-enabled compliance rather than punishment of vulnerability. Technology-mediated risk stratification delivers precise justice RED-flagged forgeries trigger investigation GREEN-verified invoices confer seamless ITC eligibility tiered YELLOW/ORANGE bands enable officer review with appeal rights, replacing arbitrary discretion with auditable reasoning.

- Three immediate policy recommendations establish implementation momentum:
- **Recommendation 1: GSTN Supplementary Verification Service.** Within 6 months,

GSTN should deploy AI verification as an optional portal service enabling MSMEs to upload purchase invoices for instant risk scoring and downloadable compliance reports prior to GSTR-3B filing. Development leverages open-source OCR/CNN frameworks with ₹5-10 crore allocation covering dataset curation and pilot testing across 10,000 invoices.

- **Recommendation 2: CBIC Due Diligence Guidance.** CBIC should issue an enforcement circular confirming that GREEN/YELLOW AI verification results constitute prima facie Section 122(2) “reasonable care” evidence, shielding MSMEs from 100% penalties when relying on official tooling. No fiscal outlay required; aligns prosecutorial discretion with technological reality.
- **Recommendation 3: Statutory Formalization.** Within 18-24 months, amend Rule 48(4) to integrate AI visual verification into IRP workflows and Rule 86A to specify risk-based ITC blocking criteria (>80% RED triggers auto-hold; <20% GREEN auto-passes). Funds existing GSTN modernization budgets; precedents include e-way bill statutory hardening.

GST Administration 3.0 envisions tax administration evolving from post-facto fraud discovery through network audits to ex-ante invoice validation at acceptance, serving dual imperatives of enforcement rigor and taxpayer protection. Fraudsters confront systematically unclaimable forgeries; MSMEs gain accessible defences against sophisticated deception; constitutional justice materializes through precision rather than presumption. As AI forgery tools proliferate and evasion tactics accelerate, policymakers confront not technical feasibility but strategic necessity: can GST survive 2026’s threat landscape on 2020-era defences? The layered verification architecture constitutes not optional modernization but constitutional imperative.