Reframing "Technological Due Process" For Tax: Adapting Administrative Law Principles to Ai-Driven Audits, Automated Assessments, and Risk Scoring

Charles Dickson Usen

Getcorp Payroll Accounting & Tax

DOI: https://doi.org/10.47772/IJRISS.2025.91200070

Received: 10 December 2025; Accepted: 19 December 2025; Published: 31 December 2025

ABSTRACT

The digital transformation of tax administration is fundamentally altering the exercise of governmental taxing authority. Rather than reviewing tax returns retrospectively, revenue authorities are increasingly monitoring taxpayers in real time through data sharing, embedded regulatory rules, and algorithmic surveillance. As the IRS advances toward "Tax Administration 3.0," supported by increased funding and artificial intelligence tools, procedural protections developed in the twentieth century are becoming increasingly inadequate.

This article examines how modern Artificial Intelligence (AI) systems, especially unsupervised machine learning and neural networks, fit within constitutional due process requirements. Its central claim is that traditional administrative law does not adequately address the challenges posed by opaque "black box" models, especially in audit selection and automated assessments. Drawing on experts' scholarship, the article proposes a new framework, "Technological Due Process 2.0," which emphasizes counterfactual explanations, system-level audits, and qualified transparency. This framework aims to uphold legitimacy in a tax system increasingly governed by algorithms rather than human discretion. The main sections are as follows: first, an exploration of AI-driven changes in tax administration; second, an analysis of the legal challenges and implications of AI models in this context; and finally, proposed solutions to bridge the gap between technology and due-process standards.

Keywords: Technological Due Process; Tax Administration; Algorithmic Audit Selection; Black Box Models; Counterfactual Explanations

INTRODUCTION

The Algorithmic Transformation of The Taxing Power

Tax administration worldwide is undergoing a profound shift. The OECD describes this as a move from "Tax Administration 2.0," which mainly digitized paperwork, to "Tax Administration 3.0," where tax rules are embedded directly into the systems people and businesses already use. In that world, compliance is built in by design, and tax authorities receive data continuously instead of waiting for annual filings.(Ruz, 2025)(OECD, 2020)

Real-time data and automation provide measurable gains for tax agencies and taxpayers—such as reducing the tax gap, compliance costs, and enforcement burdens. Yet delegating audit decisions to opaque models raises challenges to principles such as notice, explanation, and the right to be heard. While these systems are often adopted to prevent gaming and boost efficiency, these justifications may fall short of fully addressing due process concerns. As AI becomes a more prominent tool in tax administration, it is crucial to consider its potential impacts on taxpayer rights and constitutional safeguards. To maintain transparency and accountability, the adoption of explainable AI is therefore indispensable (Battaglini et al., 2024; Shaikh & Sohail, 2025; OECD, 2023).

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

In the United States, these changes are not theoretical. The Inflation Reduction Act of 2022 devoted substantial funding to modernizing IRS enforcement, and by late 2024, the Treasury Inspector General for Tax Administration had identified dozens of active AI projects, ranging from chatbots to sophisticated risk-scoring tools. (Office of Inspector General, 2024) This marks a shift from what might be called a "human-centric" regime, where civil servants exercised discretion subject to review, to an "algorithm-centric" regime in which software encodes that discretion at scale.(Siva, 2025)

This article addresses the resulting "administrative law deficit." It draws on Danielle Citron's concept of "technological due process" as applied to taxation. In this context, governmental power to seize property is increasingly mediated by opaque machine learning systems. The article argues that constitutional and administrative law must evolve to ensure both efficiency and legitimacy of the tax system in the era of AI (Daly, 2024; Citron, 2007).

II. The Machinery of the Automated Revenue State

A full understanding of due process requires examining how technology evolved from linear scoring to neural networks.

A. The Legacy of DIF and Discriminant Analysis

The IRS has long used statistical tools to flag returns for further scrutiny. These tools established a procedural baseline that is now being upended. Since the late 1960s, the agency's main enforcement tool has been the Discriminant Function System (DIF), which assigns scores to returns based on similarity to patterns of noncompliance. (Guglyuvatyy, 2025)

DIF operates on the principles of classic supervised learning. Its models are trained on "ground truth" data derived from exhaustive audits, most notably under the Taxpayer Compliance Measurement Program (TCMP) and its modern successor, the National Research Program (NRP). (Yismaw et al., 2021) In these programs, a stratified sample of taxpayers undergoes line-by-line verification, creating a labeled dataset that distinguishes compliant returns from noncompliant ones. The algorithm then identifies correlations such as a specific ratio of itemized deductions to income that reliably predict a tax deficiency. (Chan et al., 2022)

Importantly, the DIF system features a "human-in-the-loop" structure aligned with traditional safeguards. A high DIF score does not trigger an automatic audit; instead, it prompts manual review by an IRS classifier. The classifier determines whether the anomaly has a valid explanation, such as a disclosure statement or one-time event, before beginning an examination. Although the initial flag is statistical, the final decision to investigate is left to human discretion.

Furthermore, although the specific DIF formulas remain confidential under *Internal Revenue Code Section 6103(b)(2)* to prevent system "gaming," the logic is fundamentally linear. The models use weighted equations; specific variables add points to a score. An experienced tax practitioner or revenue agent can look at a high-scoring return and infer the cause. For example, a taxpayer's home office deduction may be disproportionately high compared to gross receipts. Early versions of the Discriminant Function System let professionals deduce audit flags, allowing them to advise clients on how to potentially avoid scrutiny (Camp, 2023).

This level of intelligibility holds legal significance. While the DIF system is secretive about its weights, it is not so complex as to be unintelligible. It meets the administrative law requirement that agencies articulate a rational connection between the facts and the outcomes. By contrast, shifting to unsupervised neural networks removes this connection to explainable logic, replacing linear weighting with multidimensional pattern recognition beyond human understanding. This shift to black-box AI models, such as neural networks, undermines explainability and contestability, making it harder for taxpayers to understand or challenge audit selections (Shaikh, 2025; OECD, 2023).

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

B. Neural Networks and Unsupervised Learning

While DIF automates existing audit logic, the new wave of AI tools in tax administration represents a fundamental break from it. The IRS is increasingly deploying **unsupervised learning** systems and neural networks that do not rely on pre-labeled "compliant" or "noncompliant" data. Instead of being told what to look for based on past audits, these systems ingest vast quantities of raw data to cluster taxpayers and identify anomalies and correlations that human analysts have never defined and may not even recognize. (Zheng & Penetrante, 2025) (Shaikh & Sohail, 2025) (Gans et al., 2025)

The IRS's **Large Partnership Compliance (LPC)** program shows this shift. Partnerships (Form 1065) have long been hard to audit due to their tiered structures and opacity.

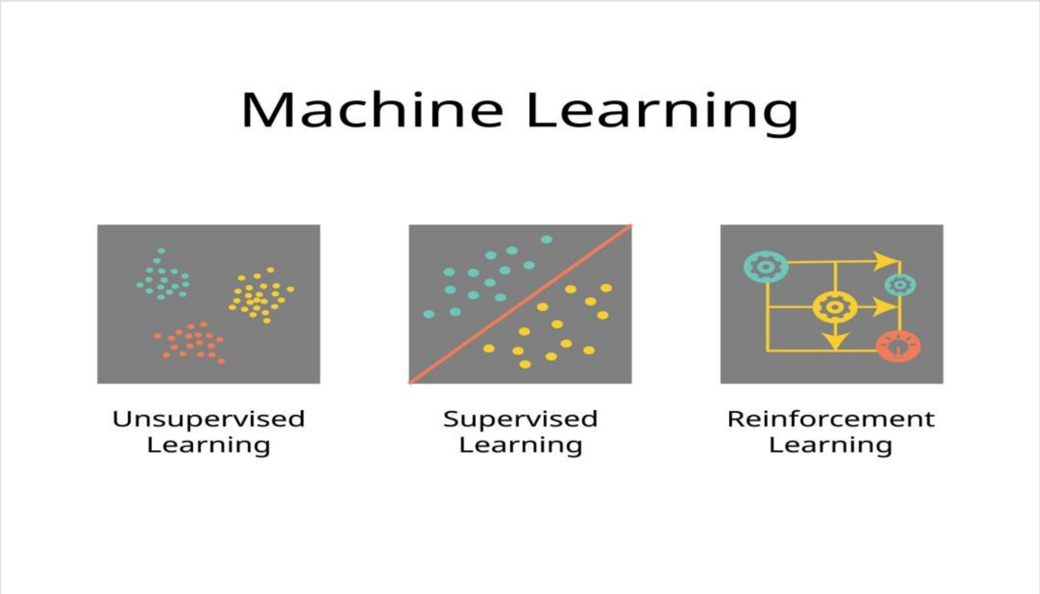

Getty Images

The LPC program uses advanced models to process not only line items but also structural relationships among partners, entities, and transactions. By mapping these complex connections, the AI flags when partnership arrangements deviate from norms learned from similar entities, rather than merely checking if a deduction exceeds a threshold. This differs from traditional DIF analysis, in which the AI identifies emergent non-compliance patterns without explicit prior programming or human-defined rules (Górski et al., 2024).

Similarly, the **Global High Wealth** program utilizes graph analytics and neural networks to build holistic, multi-entity profiles of high-net-worth individuals. Unlike DIF, which might flag a single line item, these models evaluate non-linear interactions involving hundreds or thousands of variables simultaneously. The AI might determine that a taxpayer is "high risk" not because of any single deduction, but because of a complex constellation of factors such as the timing of a transfer, the jurisdiction of a subsidiary, and the valuation method of an asset that, when taken together, statistically resemble a tax shelter. (Blank & Glogower, 2021)

The principal legal concern is the "Black Box" problem. These models utilize hidden computational layers and non-linear weighting, making it mathematically infeasible to identify a single, causal rationale for selecting a particular return. Rather than providing a clear explanation, the system effectively selects returns based on their placement within a high-risk cluster in a multidimensional space. For taxpayers seeking to contest their selection or understand the government's rationale, this opacity is complete. It eliminates the possibility of engaging in reasoned discourse with the agency, replacing the "rational connection" required by administrative law with a probabilistic determination that resists straightforward explanation (Chesterman, 2021). This fundamental shift to uninterpretable AI models poses a significant challenge to the administrative law principle of due process, depriving taxpayers of the ability to comprehend and meaningfully challenge governmental actions (Zalcewicz, 2023). Judges and agencies might consider doctrinal responses to address these challenges. One possibility is to adapt doctrines such as "arbitrary and capricious" review to require an audit trail or a simplified narrative of key

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

decision factors. Additionally, courts could mandate that agencies implementing black-box algorithms provide high-level descriptions of their algorithmic procedures to improve transparency. Such adjustments would aim to facilitate a more comprehensive judicial review and provide a framework for legal challenges, ensuring due process is maintained (OECD, 2023; Chesterman & Simon, 2021).

C. The "Unreal Audit": AUR and Math Error Authority

While unsupervised learning and neural networks capture the imagination of legal scholars, they currently target a thin slice of the tax base. For the vast majority of Americans, the automated revenue state is not a futuristic "black box" but a relentless, rules-based machine known as the Automated Underreporter (AUR) program (Camp, 2023).

Former National Taxpayer Advocate Nina Olson has famously characterized AUR as the “unreal audit.” (Olson, 2012) It functions as a massive document-matching engine, comparing the income reported on a return against the billion-plus information returns (Forms W-2, 1099, 1098) filed by third parties. When a mismatch exceeds a certain tolerance, the system automatically generates a CP2000 notice proposing additional tax. Critically, because the IRS does not legally classify this process as an "examination" under the Internal Revenue Code, taxpayers in the AUR stream are denied the statutory protections that accompany a formal audit, such as the prohibition on "unnecessary examinations" under *Section 7605(b)* or the limitations on repetitive inspections. The system effectively conducts an audit in substance while disclaiming it in form (Fogg, 2023).

Even more procedurally aggressive is the expansion of "Math and Clerical Error" authority (MEA) under IRC *Section 6213(b)*. Historically, this power was limited to obvious arithmetic mistakes, such as adding two numbers incorrectly, which allowed the IRS to summarily assess tax without issuing a Notice of Deficiency. In recent years, however, Congress and the IRS have expanded the definition of "math error" to include a wide range of "inconsistent entries" and data mismatches, particularly involving complex credits like the Earned Income Tax Credit (EITC) and the Child Tax Credit. (Guyton et al., 2018)

This procedural shift carries significant due process implications. When the IRS asserts a math error, standard deficiency procedures are suspended. Taxpayers are not afforded the customary 90-day period to petition the Tax Court; instead, they receive a notice granting only 60 days to request an abatement. If the taxpayer, often confused by the automated correspondence, fails to respond within this limited timeframe, the assessment becomes final and legally binding, resulting in the loss of the right to prepayment judicial review. Consequently, automation transforms substantive legal disputes, such as eligibility for a dependent credit, into "clerical" corrections, thereby enabling the agency to circumvent independent judicial oversight by default. (Calo & Citron, 2020)

The Erosion of Notice and Explainability

A. The Constitutional Standard

Assessing the risks associated with AI-driven enforcement requires revisiting the constitutional baseline. The Due Process Clause of the Fifth Amendment mandates that, prior to depriving an individual of property, the government must provide notice that is "reasonably calculated, under all the circumstances, to apprise interested parties of the pendency of the action." As established in *Mullane v. Central Hanover Bank & Trust Co.*, this requirement extends beyond the mere delivery of a notice; it demands substantive intelligibility. A notice that communicates the existence of a penalty but conceals its underlying rationale fails to meet this standard, as it renders the subsequent opportunity to be heard ineffective. Effective contestation is impossible if the government's reasoning remains undisclosed.

In the broader context of administrative law, this constitutional floor is reinforced by the prohibition against "arbitrary and capricious" agency action. As the Supreme Court articulated in *Motor Vehicle Mfrs. Ass’n v. State Farm Mut. Auto. Ins. Co.*, an agency must be able to articulate a "rational connection between the facts found and the choice made." This requirement presupposes a decision-making process that is legally legible, in that the agency must identify the relevant evidence and explain the logical bridge that leads to its conclusion.

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

Historically, this meant that when the IRS selected a return for audit or proposed an adjustment, a revenue agent could point to specific discrepancies such as unreported income, disallowed expenses, or misapplied credits that justified the action. The logic might have been disputed, but it was never invisible. The *reason* existed in the decision-maker's mind and could be communicated in human language. As the tax administration shifts to deep learning models, however, this assumption of explicability is collapsing. The Agency is moving from a regime of *reasoned justification* to one of *actuarial prediction*, where the "reason" for an audit is no longer a logical inference but a statistical probability derived from a high-dimensional vector space *(Hossain et al., 2025)*. (Song et al., 2025)

B. The "Black Box" Problem

If *Mullane* and *State Farm* require a "rational connection," then AI-driven tax administration presents a significant constitutional challenge. In systems governed by deep learning, the concept of meaningful notice is rapidly eroding. Taxpayers who receive generic correspondence such as "we found a discrepancy" or "your return has been selected for review" may technically receive delivery, but not substantive notice. They are not informed of the actual basis for the government's action, as this is often a high-dimensional correlation identified by a neural network that cannot be articulated by any human, including the auditor *(Greenstein, 2021)*. (OECD, 2023)

This opacity is not a monolith; it is a three-layered barrier that effectively insulates agency decisions from scrutiny:

1. **Technical Opacity:** The inherent uninterpretability of modern AI. Unlike the linear scoring of the DIF system, deep learning models utilize hidden layers of computation to identify non-linear relationships. A neural network might flag a return based on the interaction among a taxpayer's zip code, the timing of their filing, and the frequency of round-numbered expenses. While the *mathematical* path to the decision is precise, it lacks a transparent "decision tree" that can be translated into human language. The system knows *that* the return is high-risk, but it cannot explain *why* in a narrative format that satisfies due process.

2. **Legal Secrecy:** Even where models are interpretable, the IRS vigorously asserts the "law enforcement privilege" and the specific protections of I.R.C. *Section 6103(b)(2)*. These provisions allow the agency to withhold "standards used or to be used for the selection of returns" if disclosure would impair tax administration. In the analog era, this meant hiding the DIF weights. In the AI era, it means hiding the entire architecture of the surveillance state, preventing courts from examining whether the "risk" being scored is actually evidence of noncompliance or merely a proxy for poverty, race, or lawful tax avoidance.

3. **Proprietary Secrecy:** The problem is compounded when agencies procure AI tools from private vendors. In these cases, the algorithms are often shielded not just by government privilege but by intellectual property law. Vendors argue that their risk-scoring models are "trade secrets," effectively privatizing the logic of public enforcement. This creates a scenario where a taxpayer's liability may hinge on a proprietary formula that the government itself may not fully understand or own *(Christoph, 2023)*.

Current attempts to mitigate this through "Explainable AI" (XAI), such as generating "feature importance" maps that highlight which variables carried the most weight, are legally insufficient. A chart showing that "Income" or "Business Expenses" was a significant factor describes the model's *general behavior*, but it does not isolate the *dispositive reason* for a specific decision. It tells the taxpayer which variables correlated with the outcome, but it fails to answer the dispositive due process question: *What distinct fact must I challenge to reverse this decision?* Without that answer, the right to be heard is reduced to a guessing game *(Lipman, 2021)*. (Barocas et al., 2019)

The Crisis of the Hearing: Automated Adjudication and Bias

A. From Adjudication to Data Entry

As the "*notice*" component of due process deteriorates, the "*opportunity to be heard*" is experiencing an even more pronounced decline. Traditionally, administrative hearings, even informal ones, provided a forum for

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

dialogue, allowing individuals to present their case to a human decision-maker capable of evaluating evidence, assessing credibility, and exercising discretion based on the totality of circumstances. The increasing use of AI in tax administration undermines this essential aspect of procedural fairness, as human judgment is increasingly replaced by algorithmic determinations *(Guglyuvatyy, 2025). (OECD, 2023)*

In the context of automated revenue administration, the dialogic process is increasingly reduced to a unidirectional data-entry exercise. For millions of taxpayers, the modern "hearing" is no longer a meeting or telephone conversation, but rather a correspondence interaction processed by optical character recognition (OCR) and inflexible workflow software. When taxpayers respond to a CP2000 underreporter notice or a math error assessment, their correspondence is frequently processed by systems that search for specific keywords, numerical values, or standardized forms.

This technological filter creates a profound procedural bottleneck. If a taxpayer's defense relies on nuance, for example, arguing that a missed deadline was due to "reasonable cause" arising from a medical emergency or a natural disaster, the automated system is often blind to it. Unless the explanation fits into a pre-defined data field or box, it effectively does not exist *(Grant et al., 2023). (OECD, 2023) (OECD, 2023)*

The human agents who do review these responses, often termed "screen-level bureaucrats," are frequently constrained by the same rigid software. They may lack the authority (or the time) to override the system's logic or to consider evidence that falls outside the standard operating procedure. The result is a system of "adjudication by algorithm," where the complexity of human life is flattened into binary code. A letter explaining a good-faith error is treated simply as "Response Received: Invalid Format," effectively denying the taxpayer a meaningful opportunity to be heard before the liability is finalized *(Johnson, 2015)*. This shift transforms the administrative review process into a perfunctory exercise, where algorithmic efficiency supersedes substantive due process and the nuanced consideration of individual circumstances *(Calo & Citron, 2020)*.

The degradation of the human hearing is particularly dangerous because it removes the primary safety valve for detecting and correcting systemic bias. In a manual system, a pattern of inequitable enforcement might eventually be noticed by agents or supervisors. In an automated system, however, bias can be codified and scaled efficiently, hidden behind the veneer of mathematical neutrality.

A critical failure of the automated revenue state is its susceptibility to "proxy discrimination." This phenomenon was starkly illustrated by a comprehensive 2023 empirical study conducted by researchers at Stanford University and the U.S. Treasury. The study revealed that the IRS audits Black taxpayers claiming the Earned Income Tax Credit (EITC) at a rate between **2.9 and 4.7 times higher** than non-Black taxpayers claiming the same credit *(Miller, 2023). (Elzayn et al., 2023).*

Crucially, this disparity did not arise because the IRS programmed its computers to target Black taxpayers, as the agency's algorithms are legally forbidden from using race as a variable. Instead, the disparity emerged because the algorithms were optimizing for "audit efficiency," which is defined as minimizing the administrative cost per dollar of revenue protected. *(Elzayn et al., 2023)* This metric, while seemingly neutral, inadvertently serves as a proxy for race when combined with other demographic and socioeconomic factors that are frequently correlated with the EITC recipient population *(Black et al., 2022). (Hertz & Tom, 2023)*

The EITC is a complex credit often claimed by low-income filers with fluctuating family structures. Auditing these claims is "cheap" for the agency: it can be done entirely through automated correspondence (the "unreal audit") without deploying expensive field agents. By contrast, auditing a high-net-worth partnership requires specialized human talent and hundreds of hours. Consequently, an unsupervised learning model instructed to "maximize the number of noncompliant returns identified per budget dollar" will naturally learn to target EITC claimants. *(Henderson et al., 2023)* (Elzayn et al., 2023)

Because race in the United States is statistically correlated with certain economic and demographic markers such as zip code, family size, and single-parent household status, the algorithm effectively found "proxies" for race. It learned that targeting these specific demographic clusters yielded the highest volume of easy "wins." The result is a form of structural discrimination that is invisible to the Equal Protection Clause's requirement of "discriminatory intent," yet devastating in its impact. The machine is not "racist" in a human sense as it is simply ruthlessly efficient at targeting the vulnerable, creating a feedback loop where the poor are over-policed because they are cheaper to police *(Mayson, 2019)*. (Mayson & G., 2019)

C. The *Mathews* Calculus

The constitutional validity of these automated procedures ultimately rests on the balancing test established in *Mathews v. Eldridge*. This doctrine dictates that the specific dictates of due process are determined by weighing three distinct factors: (1) the private interest affected by the official action; (2) the risk of an erroneous deprivation of such interest through the procedures used, and the probable value, if any, of additional or substitute procedural safeguards; and (3) the government's interest, including the function involved and the fiscal and administrative burdens that the additional or substitute procedural requirement would entail.

For decades, the *Mathews* calculus has heavily favored the IRS. Courts have generally accepted that the sheer volume of tax returns makes individualized hearings or detailed, custom-written explanations for every audit selection prohibitively expensive (Factor 3). The administrative burden of "human-centric" due process was simply too high.

The advent of AI, however, fundamentally alters the variables of this equation, specifically the third factor. In the analog era, providing a "reasoned explanation" required a human agent to sit down, review a file, and dictate a letter, a costly process that scales linearly with the number of audits. In the AI era, the marginal cost of generating an explanation is plummeting toward zero (*Kesari et al., 2024*). Advanced natural language generation models can now synthesize coherent, contextually relevant explanations for algorithmic decisions with minimal computational overhead, effectively reducing the administrative burden associated with robust procedural safeguards (*Coglianese, 2021*). For example, early pilot programs in the public sector have demonstrated this potential: the UK Government's Office of Communications has used AI-generated explanations to efficiently handle compliance audits, achieving faster, clearer outcomes. Similarly, a study conducted in collaboration with the German Federal Employment Agency found that AI-driven explanations in social program adjudications significantly improved participants' understanding and reduced formal complaints.

The same technologies that power automated enforcement, such as Large Language Models (LLMs) and automated reasoning systems, can be configured to generate standardized, plain-language explanations of *why* a return was flagged. If an algorithm can calculate a risk score in milliseconds, it can arguably also output the primary factors driving that score (or a counterfactual explanation) just as instantly.

Consequently, the government's traditional argument that providing more processes is "too burdensome" is losing its factual predicate. We are entering a period where the risk of erroneous deprivation (Factor 2) is rising due to the opacity and bias of "black box" models, while the cost of providing additional safeguards (Factor 3) is falling due to the efficiency of those same models. Under a faithful application of *Mathews*, this shift demands a recalibration: as the cost of explanation drops, the constitutional floor for what counts as "sufficient process" must rise. The refusal to explain is no longer an economic necessity; it is a policy choice (Mei & Broyde, 2025; Vredenburgh, 2021). (Mei et al., 2025)

Judicial Review and the Wall of Secrecy

Even if a taxpayer suspects that they have been targeted by a biased or malfunctioning algorithm, they face a formidable jurisdictional barrier when seeking relief in federal court. Under the *Administrative Procedure Act (APA) Section 701(a)(2)* and the Supreme Court's ruling in *Heckler v. Chaney*, agency decisions regarding whether to investigate, prosecute, or enforce are presumptively "committed to agency discretion by law."

This doctrine rests on the practical theory that courts are ill-equipped to second-guess how agencies allocate their scarce resources. Judges, the logic goes, cannot determine whether an agency should spend its budget auditing a partnership in New York versus a sole proprietor in Texas. In the tax context, this has solidified into a near-total immunity for audit-selection decisions. When a taxpayer challenges a deficiency in Tax Court, the

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

scope of review is typically limited to the *correctness* of the tax assessment (de novo review), not the *propriety* of the investigation that led to it. The "why" of the audit is treated as legally irrelevant.

In an age of algorithmic selection, however, this hands-off posture carries new and profound dangers. The *Heckler* doctrine was designed to protect the ad-hoc, case-by-case judgment calls of human prosecutors. It was not designed to shield a fully automated, programmatic system that screens millions of citizens simultaneously.

When an AI model selects targets based on proxy variables for race or class (as discussed in Section IV), it is not making an individual resource-allocation decision; it is executing a systemic enforcement policy. Yet, under current doctrine, if the resulting tax assessment is technically correct, that is, the taxpayer actually owes the money, the court will almost invariably uphold the result. The judiciary's focus on the *individual outcome* blinds it to systemic *"harm"*. A taxpayer can be "legally" taxed even if they were "illegally" selected, meaning that discriminatory algorithms can operate indefinitely without ever triggering a reversible error in court. *(Von Ungern-Sternberg, 2022)*

B. The *Armstrong* Barrier

If the "committed to agency discretion" doctrine is the shield that protects algorithmic enforcement, *United States v. Armstrong* is the sword that strikes down attempts to challenge it. In that seminal case, the Supreme Court held that a defendant seeking discovery to support a selective prosecution claim must first make a "credible showing" of distinct discriminatory effect *and* discriminatory intent.

In the context of AI-driven tax administration, this requirement creates a lethal circularity, a "Catch-22" of constitutional proportions. To survive a motion to dismiss, a taxpayer must allege specific facts showing that the IRS's algorithm was designed with discriminatory intent or is producing a disparate impact. However, the evidence needed to prove those facts, such as the algorithm's weighting factors, training data, or internal validation reports, is in the government's exclusive control and is shielded by the "law enforcement privilege" and trade secrecy protections discussed earlier.

Under *Armstrong*, the taxpayer cannot get discovery (access to the model) without first showing evidence of bias. But they cannot show evidence of bias without first getting discovery. In the analog world, a plaintiff might overcome this by finding a whistleblower or a leaked memo. In the digital world, where the "intent" of the system is buried in millions of parameters of a neural network, such smoking guns rarely exist.

This doctrinal loop effectively immunizes algorithmic bias from constitutional scrutiny. Unless the judiciary adjusts the *Armstrong* framework, for instance, by allowing limited, supervised access to models for independent experts upon a statistical showing of disparity, taxpayers are left with a right that has no remedy. The government can deploy systems that disproportionately audit protected groups, and so long as the discrimination is "unintentional" and the model is kept secret, the courthouse doors remain locked.

Comparative Lessons: *SyRI* and the GDPR

A. The SyRI Precedent: A Judicial Rebuke of the Black Box

While U.S. courts grapple with applying 20th-century administrative law to 21st-century technology, the Dutch judiciary has provided a strikingly relevant roadmap for the future. In the landmark 2020 *SyRI* (System Risk Indication) judgment, the District Court of The Hague struck down a government program that used algorithmic risk scoring to identify potential welfare fraud. (Netherlands: Court Prohibits Government's Use of AI Software to Detect Welfare Fraud, 2020) (Appelman et al., 2021)

The parallels to the IRS's modernization efforts are exact: SyRI aggregated data from multiple government silos, such as taxes, employment, immigration, and benefits, to generate risk profiles for individual citizens. The Dutch court held that this system violated Article 8 of the European Convention on Human Rights (the right to private life) not because the goal of fraud detection was illegitimate, but because the *method* was opaque. The court explicitly condemned the "black box" nature of the system, noting that because the risk model's indicators and

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

weights were secret, citizens could not understand why they were targeted or effectively defend themselves against the accusation of "high risk."(Usen, 2025)

Crucially, the judgment established a principle that U.S. courts have been reluctant to embrace: that **transparency is a precondition for legality** in automated enforcement. The court reasoned that without *verifiability*, the ability for an independent party to check the algorithm's logic, there is no check on power. A system that cannot be scrutinized is, by definition, disproportionate in a democratic society. This decision underscores the fundamental rights to an effective remedy and to due process, highlighting the imperative of transparent algorithmic decision-making, particularly in contexts that impact fundamental rights *(Kuźniacki et al., 2022; Papis-Almansa et al., 2022)*. (Yalcin et al., 2022).

B. The GDPR and the Right to Explanation

Beyond case law, the European Union has codified these protections in the General Data Protection Regulation (GDPR). Unlike the U.S. Privacy Act of 1974, which was built for file cabinets, the GDPR was built for data streams. Articles 13, 14, and 15 grant individuals the right to access not only their raw data but also "meaningful information about the logic involved" in any automated decision-making process. (Wachter et al., 2017)

Furthermore, Article 22 creates a presumptive right not to be subject to a decision based solely on automated processing if it produces legal effects. While there are exceptions for tax and fraud prevention, the core *right to explanation* remains a powerful normative baseline. It compels agencies to design systems that are interpretable by design, ensuring that when a decision is made, a human-readable explanation can be generated *(Frész et al., 2024).*

In stark contrast, the United States suffers from a "statutory void." There is no federal equivalent to Article 22 governing the IRS. Taxpayers are left to rely on the *Administrative Procedure Act* (1946) and the *Due Process Clause* (1791), neither of which was written to contemplate neural networks. This comparative silence highlights that the "administrative law deficit" in U.S. tax administration is not an inevitable consequence of technology, but a specific failure of legislative will (Guglyuvatyy, 2025; Usen, 2025).

Technological Due Process 2.0 aims to build into U.S. tax law what courts and regulators in Europe are already demanding for high-risk AI systems: real transparency, serious impact assessment, and usable explanations as prerequisites for deploying those systems at all. These safeguards are not framed as a new, standalone data-protection regime, but instead are enforced through familiar U.S. administrative-law tools like the Mathews due process test, the State Farm "arbitrary and capricious" standard, and traditional Administrative Procedure Act record review.

Towards Technological Due Process 2.0

To address the tension between the automated revenue state and the rule of law, this article proposes a reformed framework.

A. Doctrinal Roadmap: From Framework to Enforceable Standards

"Technological Due Process 2.0" is not offered as a freestanding constitutional innovation. Rather, it is a way of cashing out existing public-law commitments in an automated revenue state. In concrete terms, the framework translates into three clusters of enforceable legal standards.

First, arbitrary and capricious review under *State Farm* becomes a demand for reasoned algorithmic decision-making. When AI tools drive audit selection or automated assessments, agencies should be required to preserve a reviewable audit trail such as counterfactual explanations and model-level documentation that shows a rational connection between the input data and the outcome in the administrative record.

Second, the *Mathews v. Eldridge* balancing test supplies the baseline for calibrating procedure. As Parts III and IV show, black-box models increase the risk of erroneous deprivation, while modern language models dramatically lower the marginal cost of individualized explanations and additional safeguards. Under *Mathews*,

ISSN No. 2454-6186 | DOI: 10.47772/IJRISS | Volume IX Issue XII December 2025

those shifts in factors two and three warrant more robust notice, explanation, and human review in precisely the high-risk settings where AI is being deployed.

Third, the APA's record-review and transparency norms extend naturally to algorithmic systems. On this view, training-data documentation, validation studies, bias-testing results, and explanation templates are all part of the "whole record" that must be available, at least in camera, when courts review AI-mediated tax enforcement. Technological Due Process 2.0 therefore asks courts to treat algorithmic models not as inscrutable enforcement black boxes, but as agency reasoning that must be open to ordinary APA scrutiny.

B. Reframing Notice: Counterfactual Explanations

The first pillar of "Technological Due Process 2.0" necessitates a fundamental redefinition of *notice*. In the context of black-box algorithms, notifications that merely state the outcome ("Audit Selected") are constitutionally insufficient, as they provide no basis for contestation. Conversely, requiring full transparency, such as publishing source code or neural network weights, is often legally precluded by trade secrecy and, in practice, ineffective due to technical complexity.

A more effective approach involves shifting the standard from **transparency** (access to code) to **interpretability** (understanding system behavior). The most promising mechanism for achieving this objective is the **Counterfactual Explanation**, a concept introduced in the context of the GDPR by legal scholar Sandra Wachter (Wachter et al., 2017). (Wachter et al., 2017)

A counterfactual explanation does not attempt to describe the internal logic of the "black box." Instead, it describes the decision's external boundary. It answers the specific question: *"What is the minimal change in the input data that would have produced a different outcome?"*

For example, consider a taxpayer selected for an audit by a neural network analyzing Schedule C returns. A traditional explanation might be silence, or a confusing "feature importance" map showing that "Expenses" were weighted heavily. A counterfactual explanation, however, would state:

"You were flagged for audit because your ratio of Business Expenses to Gross Receipts is 65%. If this ratio had been 40% or lower, you would not have been selected."

This approach offers a decisive legal breakthrough for three reasons:

1. **Actionability:** It restores the "opportunity to be heard" by providing the taxpayer with a concrete, factual target. The taxpayer now knows they do not need to challenge the algorithm's math and that they must prove their 65% expense ratio is legitimate and supported by evidence.

2. **Protection of Government Interests:** It preserves the "law enforcement privilege." The IRS does not have to reveal the entire risk model, the specific weights for other variables, or the vendor's proprietary logic. It only reveals the specific threshold that was crossed in this specific case.

3. **Prospective Compliance:** This approach transforms enforcement into an educational process. By identifying the "tipping point," taxpayers and practitioners gain insight into the boundaries of compliant behavior. Instead of speculating about hidden rules, they receive actionable feedback that informs future reporting, thereby supporting the overarching objective of voluntary compliance (World Bank, 2019). Although counterfactual explanations enhance individual procedural justice, they are structurally inadequate for identifying or remedying systemic bias. As discussed in Section V, the *Armstrong* barrier effectively precludes individual plaintiffs from establishing discrimination, as they lack access to the aggregate data required to demonstrate disparate impact. An individual taxpayer is aware only of their own outcome and cannot ascertain whether the algorithm is similarly flagging a disproportionate number of others.o flagging 90% of their neighbors.

Therefore, the remedy for algorithmic harm must be structural and *ex ante*. We cannot wait for damage to occur; we must prevent it. The IRS must implement rigorous **Algorithmic Impact Assessments (AIA)** before deploying any high-stakes enforcement model.

This recommendation aligns with the emerging federal standard set by the Office of Management and Budget in **Memo M-24-10**. This guidance directs federal agencies to assess AI systems for "risks to rights and safety" before they are operationalized. In the specific context of tax administration, an AIA would function as a "stress test" similar to those used in banking regulation. (*Office of Management and Budget, 2024*) Before a new audit-selection model is turned loose on the American public, it must be run against historical data to answer critical questions:

- Does this model produce more false positives for protected classes?
- Does it use proxies (like zip code or family structure) that effectively reconstruct race?
- Does the "efficiency" gain justify the disparate impact?

It is essential that these assessments are not conducted solely as internal exercises. The IRS should not serve as the exclusive evaluator of its own systems. Instead, such assessments should be subject to review by an independent oversight body with the technical expertise to "audit the auditors." The **National Taxpayer Advocate (NTA)** is well-positioned to fulfill this role.

The NTA already possesses a statutory mandate under *I.R.C. Section 7803(c)* to identify systemic problems in tax administration and report them to Congress. This mandate should be updated to include the "constitutionality code" of the IRS's AI. If a proposed algorithm shows a significant bias against low-income EITC claimants, as identified in the 2023 Stanford study, the NTA should have the power to flag this risk and halt deployment until the agency can demonstrate that the disparity is necessary to achieve a compelling government interest, rather than merely a byproduct of cost-cutting. *(Cords, 2021; Lipman, 2021) (Elzayn et al., 2023).*

C. Qualified Transparency

The final necessary reform lies within the judiciary itself. To break the deadlock created by *United States v. Armstrong*, where plaintiffs are denied discovery because they lack the evidence that only discovery can provide, courts must adopt a regime of **Qualified Transparency**.

Full public disclosure of IRS algorithms is likely impossible; the agency has legitimate interests in protecting "law enforcement logic" from being gamed by evaders, and private vendors have legitimate intellectual property rights in their code. However, "secrecy from the public" need not equate to "secrecy from the court."

Courts should utilize protective orders and **in camera review** to permit limited access to contested algorithms. Under this framework, the proprietary "black box" is not opened to the world, but to a specific set of eyes: **court-appointed technical experts** or Special Masters (under *Federal Rule of Civil Procedure 53* or *Rule 706*). These cleared experts can audit the model, training data, and weighting factors under strict confidentiality to determine whether the "risk" was calculated using impermissible proxies such as race or religion. This effectively balances the government's need for security with the citizen's right to a fair trial.

Furthermore, courts must stop treating AI outputs as incontestable facts and start treating them as expert opinions subject to the **Daubert standard**. In *Daubert v. Merrell Dow Pharmaceuticals*, the Supreme Court designated trial judges as "gatekeepers" responsible for ensuring that scientific testimony is not only relevant but reliable.

An automated risk score or an AI-generated assessment is, in essence, a form of expert testimony, an opinion derived from a complex methodology. Therefore, before the government can introduce an AI's conclusion as evidence of a deficiency, it must lay a foundation for the model's reliability. It must show:

1. Has the model been tested?

2. What is its known or potential error rate?
3. Has it been subjected to peer review or independent validation?

If the IRS cannot answer these questions because the model is a "black box," the evidence should be inadmissible. Applying *Daubert* forces the agency to "show its work," ensuring that no citizen loses property to "junk science" hidden in a computer.

D. Tiered Transparency Obligations in Tax Administration

Experience with rules like the GDPR and the draft EU Artificial Intelligence Act shows that transparency and oversight work best when they are matched to the level of risk, not imposed as a blunt, all-or-nothing rule. Taking that lesson into tax administration, Technological Due Process 2.0 lays out a three-tier framework.

For low-impact tools, such as internal analytics used only inside the tax agency, the system would only require broad public explanations of what the tools do and how they are used. For medium-impact tools, like automated correspondence systems (including AUR mismatch notices and expanded math-error streams), agencies would have to provide standard notice templates and personalized “what-if” explanations that show how different facts or inputs would change the outcome, before a tax bill is locked in, with particular attention to low-income taxpayers and EITC recipients.

For high-impact tools, especially complex neural-network models used in programs like Large Partnership Compliance and Global High Wealth, the rules would be much stricter: external Algorithmic Impact Assessments, independent bias testing, and confidential (in camera) review of the models themselves, coupled with detailed, taxpayer-specific explanations and a clear right to human review of automated decisions.

CONCLUSION

The transition to "Tax Administration 3.0" is not merely a technical upgrade. It is an inevitable structural evolution of the state, driven by the fiscal imperative to close the tax gap and the technological capability to monitor economic activity in real time. The promise of this shift is undeniable: a world of seamless compliance, reduced friction, and optimized revenue collection. However, as Justice Holmes in *Compania General De Tabacos De Filipinas v. Collector of Internal Revenue*, 275 U.S. 87, 100 (1927) famously noted, "Taxes are what we pay for civilized society," and a society cannot call itself civilized if the mechanisms of its funding are inscrutable to its citizens.

Efficiency cannot be purchased at the price of legitimacy. A tax system that operates as a "black box" where liability is determined by opaque algorithms, enforced by automated notices, and insulated from judicial review will ultimately erode the foundation of the American revenue system: voluntary compliance. If taxpayers come to view the IRS not as a fair administrator of law but as an arbitrary, invincible machine, the social contract that undergirds the tax system will fracture.

The appropriate path is not to reject automation, but to ensure its constitutional integration. By adopting the framework of "Technological Due Process 2.0," incorporating **counterfactual explanations** to restore notice, **system-level auditing** to promote equity, and **qualified transparency** to facilitate judicial oversight, the United States can adapt foundational principles of the Rule of Law to the contemporary reality of algorithmic governance. Specific legislative actions could include mandating algorithmic impact assessments for all AI tools used in tax administration and ensuring they do not disproportionately affect protected classes. Furthermore, regulatory bodies could be tasked with establishing guidelines requiring transparency and interpretability standards for AI deployments. At the judicial level, courts could adopt a standard treating AI determinations as expert testimony, requiring rigorous scrutiny under the Daubert standard as a prerequisite for the admissibility of AI-generated audit selections. As the tax code becomes code, it must remain subject to the fundamental "source code" of democracy: the Due Process Clause.

REFERENCES

1. Appelman, N., Fahy, R., & Hoboken, J. V. (2021). Social Welfare, Risk Profiling and Fundamental Rights: The Case of SyRI in the Netherlands. JIPITEC, 12. https://www.ivir.nl/publications/social-welfare-risk-profiling-and-fundamental-rights-the-case-of-syri-in-the-netherlands/
2. Barocas, S., Selbst, A. D., & Raghavan, M. (2019). The Hidden Assumptions Behind Counterfactual Explanations and Principal Reasons. arXiv. https://arxiv.org/abs/1912.04930
3. Battaglini, M., Guiso, L., Lacava, C., Miller, D. L., & Patacchini, E. (2024). Refining public policies with machine learning: The case of tax auditing. Journal of Econometrics, 249, 105847. https://doi.org/10.1016/j.jeconom.2024.105847
4. Black, E., Elzayn, H., Chouldechova, A., Goldin, J., & Ho, D. (2022). Algorithmic Fairness and Vertical Equity: Income Fairness with IRS Tax Audit Models. 2022 ACM Conference on Fairness, Accountability, and Transparency, 1479–1503. https://doi.org/10.1145/3531146.3533204
5. Blank, J. D., & Glogower, A. (2021). The trouble with targeting tax shelters. Administrative Law Review. https://ssrn.com/abstract=3937450
6. Caliendo, P. (2025). Algorithmic Tax Power and Constitutional Safeguards: Global Perspectives on AI, bias, and Digital Tax justice. Beijing Law Review, 16(03), 1861–1879. https://doi.org/10.4236/blr.2025.163094
7. Calijuri, M. S. (2023). Revolutionizing Taxation: Tax Administration 3.0 is already here. Review of International and European Economic Law, 2(4), A4.1-A4.14. https://rieel.com/index.php/rieel/article/view/78
8. Calo, R., & Citron, D. K. (2020). The Automated Administrative State: A Crisis of Legitimacy. Emory Law Journal, 70. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3553590
9. Camp, B. T. (2023). Collective due process in tax administration. Pittsburgh Tax Review, 20(1). https://doi.org/10.5195/taxreview.2022.178
10. Chan, T., Tan, C., & Tagkopoulos, I. (2022). Audit lead selection and yield prediction from historical tax data using artificial neural networks. PLoS ONE, 17(11), e0278121. https://doi.org/10.1371/journal.pone.0278121
11. Chesterman, S. (2021). Through a glass, darkly: Artificial intelligence and the problem of opacity. The American Journal of Comparative Law, 69(2), 271–294. https://doi.org/10.1093/ajcl/avab012
12. Christoph, M. (2023). Criminal Justice Technology and the Regulatory Sandbox: Toward balancing justice, accountability, and innovation. University of Pittsburgh Law Review, 84(4). https://doi.org/10.5195/lawreview.2023.959
13. Citron, D. K. (2007). Technological due process. Washington University Law Review, 85(6), 1249–1313. https://digitalcommons.law.umaryland.edu/fac_pubs/406
14. Coglianese, C. (2021). Administrative law in the automated state. Daedalus, 150(3), 104–120. https://doi.org/10.1162/daed_a_01862
15. Cords, D. W. (2021). Nina E. Olson: A Legislative legacy. Pittsburgh Tax Review, 18(1). https://doi.org/10.5195/taxreview.2020.126
16. Daly, S. (2024). Artificial Intelligence, the Rule of Law and Public Administration: The Case of Taxation. The Cambridge Law Journal, 83(3), 437–464. https://doi.org/10.1017/s0008197324000448
17. Daubert v. Merrell Dow Pharmaceuticals, Inc., 509 U.S. 579 (1993)
18. Elzayn, H., Smith, E., Hertz, T., Guage, C., Ramesh, A., Fisher, R., Ho, D. E., & Goldin, J. (2024). Measuring and mitigating racial disparities in tax audits. The Quarterly Journal of Economics, 140(1), 113–163. https://doi.org/10.1093/qje/qjae027
19. Fogg, K. (2023). Revisiting the ten deadly sins created in the IRS Restructuring and Reform Act. Pittsburgh Tax Review, 20(1). https://doi.org/10.5195/taxreview.2022.176
20. Fresz, B., Dubovitskaya, E., Brajovic, D., Huber, M., & Horz, C. (2024). How should AI decisions be explained? Requirements for Explanations from the Perspective of European Law. arXiv. https://doi.org/10.48550/arxiv.2404.12762
21. Gans, J., Logan, P., & Piper, M. A. (2025). IRS Audits and the Emerging Role of AI in Enforcement. Holland & Knight Insights. https://www.hklaw.com/en/insights/publications/2025/11/irs-audits-and-the-emerging-role-of-ai-in-enforcement

22. Górski, Ł., Kuźniacki, B., Almada, M., Tyliński, K., Calvo, M., Asnaghi, P. M., Almada, L., Iñiguez, H., Rubianes, F., Pera, O., & Nigrelli, J. I. (2024). Exploring explainable AI in the tax domain. Artificial Intelligence and Law, 33(3), 551–579. https://doi.org/10.1007/s10506-024-09395-w
23. Grant, D. G., Behrends, J., & Basl, J. (2023). What we owe to decision-subjects: beyond transparency and explanation in automated decision-making. Philosophical Studies, 182(1), 55–85. https://doi.org/10.1007/s11098-023-02013-6
24. Greenstein, S. (2021). Preserving the rule of law in the era of artificial intelligence (AI). Artificial Intelligence and Law, 30(3), 291–323. https://doi.org/10.1007/s10506-021-09294-4
25. Guglyuvatyy, E. (2025). Balancing innovation and integrity: AI in tax administration and taxpayer rights. Humanities and Social Sciences Communications, 12(1). https://doi.org/10.1057/s41599-025-06099-7
26. Guyton, J., Leibel, K., Manoli, D., Patel, A., Payne, M., & Schafer, B. (2018). The effects of EITC Correspondence Audits on Low-Income Earners. NBER Working Paper No. 24465. https://doi.org/10.3386/w24465
27. Heckler v. Chaney, 470 U.S. 821 (1985)
28. Henderson, P., Chugg, B., Anderson, B., Altenburger, K., Turk, A., Guyton, J., Goldin, J., & Ho, D. E. (2023). Integrating reward maximization and population estimation: Sequential Decision-Making for Internal Revenue Service Audit selection. Proceedings of the AAAI Conference on Artificial Intelligence, 37(4), 5087–5095. https://doi.org/10.1609/aaai.v37i4.25637
29. Hossain, N. I., Ahmad, N. R., Amin, N. M. R., Sultana, N. N., & Chowdhury, N. S. (2025). Risk assessment in financial services: Advancing transparency and trust through explainable AI models. International Journal of Science and Research Archive, 16(1), 1409–1419. https://doi.org/10.30574/ijsra.2025.16.1.2150
30. Inflation Reduction Act of 2022, Pub. L. No. 117-169, 136 Stat. 1818.
31. Internal Revenue Manual (IRM) § 4.1.
32. Internal Revenue Service, Strategic Operating Plan FY 2023-2031.
33. Johnson, S. R. (2015). Reasoned explanation and IRS adjudication. Duke Law Journal, 63(8). https://scholarship.law.duke.edu/dlj/vol63/iss8/5
34. Kesari, A., Sele, D., Ash, E., & Bechtold, S. (2025, May 6). Explaining eXplainable AI. SSRN. https://doi.org/10.2139/ssrn.4972085
35. Kuzniacki, B., Almada, M., Tyliński, K., & Górski, Ł. (2022). Requirements for tax XAI under constitutional principles and human rights. In Lecture notes in computer science (pp. 221–238). https://doi.org/10.1007/978-3-031-15565-9_14
36. Lipman, F. J. (2021). Humanizing the tax system: What national taxpayer advocate Nina E. Olson did for America's kids and their families. Pittsburgh Tax Review, 18(1). https://doi.org/10.5195/taxreview.2020.121
37. Maria Amparo Grau Ruiz, Fiscal Transformations Due to AI and Robotization: Where Do Recent Changes in Tax Administrations, Procedures and Legal Systems Lead Us?, 19 Nw. J. Tech. & Intell. Prop. 325. https://scholarlycommons.law.northwestern.edu/njtip/vol19/iss4/1
38. Mathews v. Eldridge, 424 U.S. 319 (1976)
39. Mayson, S. G. (2019, June). "Bias in, bias out." Yale Law Journal, 128, 2218–2300.
40. Mei, Y., & Broyde, M. J. (2025). Reclaiming constitutional authority of algorithmic power. arXiv. https://doi.org/10.48550/arxiv.2508.11699
41. Miller, K. (2023, January 31). IRS disproportionately audits Black taxpayers. Stanford Law School. https://law.stanford.edu/press/irs-disproportionately-audits-black-taxpayers/
42. Motor Vehicle Mfrs. Ass'n v. State Farm Mut. Auto. Ins. Co., 463 U.S. 29 (1983)
43. Mullane v. Central Hanover Bank & Trust Co., 339 U.S. 306 (1950)
44. OECD. (2020). Tax Administration 3.0: The Digital Transformation of Tax Administration. OECD. https://doi.org/10.1787/ca274cc5-en
45. OECD. (2023). AI in tax administration: Governing with Artificial Intelligence. OECD. https://www.oecd.org/en/publications/governing-with-artificial-intelligence_795de142-en/full-report/ai-in-tax-administration_30724e43.html
46. OECD. (2023). Tax Administration 2023: Comparative Information on OECD and Other Advanced and Emerging Economies. OECD.

https://www.oecd.org/content/dam/oecd/en/publications/reports/2023/09/tax-administration-2023_87655bc9/900b6382-en.pdf
47. Office of Management and Budget. (2024). Advancing governance, innovation, and risk management for agency use of artificial intelligence (Memo M-24-10). Executive Office of the President.
48. Olson, N. E. (2012). What's an Audit, Anyway? Taxpayer Advocate Service. https://www.taxpayeradvocate.irs.gov/news/nta-blog/ntablog-whats-an-audit-anyway/2012/01/
49. Olson, N. E. (2012, April 19). Written statement of Nina E. Olson, National Taxpayer Advocate... U.S. House of Representatives. https://www.taxpayeradvocate.irs.gov/wp-content/uploads/2020/09/testimony_house_oversight_comm_4-19-2012.pdf
50. Papis-Almansa, M. (2022). The use of new technologies in VAT and taxpayers' rights. In CJEU: Recent developments in value added tax 2021. Linde. https://ssrn.com/abstract=4034858
51. Regulation (EU) 2016/679 (GDPR), Arts. 13-15, 22.
52. Ruz, C. (2025, October 27). Systems integration in Tax Administration 3.0 - Part 3. CIAT. https://www.ciat.org/systems-integration-in-tax-administration-3-0/?lang=en
53. Shaikh, S. (2025). Balancing Innovation and oversight: AI in the U.S. Treasury and IRS: a survey. arXiv. https://doi.org/10.48550/arxiv.2509.16294
54. Siva, P. K. (2025). Legal and regulatory implications of Agentic AI in tax administration. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5333340
55. Song, Y., Wang, Z., Li, N., & Hongji, X. (2025). A Hybrid DNN-Transformer-AE Framework for Corporate Tax Risk Supervision and Risk Level Assessment. Economics of Management and Innovation. https://www.gbspress.com/index.php/EMI/article/view/456
56. SyRI Judgment, District Court of The Hague (C/09/550982 / HA ZA 18-541), Feb. 5, 2020.
57. United States v. Armstrong, 517 U.S. 456 (1996).
58. Usen, C. D. (2025). Analyzing behavioral economics in enhancing voluntary tax compliance in the us among self-employed informal sector workers. World Journal of Advanced Research and Reviews, 26(3), 366–384. https://doi.org/10.30574/wjarr.2025.26.3.2226
59. Usen, C. D. (2025). From Hess to O'Donnabhain: The divergence of Business utility and medical necessity in Somatic Tax Law. International Journal of Research Publication and Reviews, 6(11), 10014–10020. https://doi.org/10.55248/gengpi.06.1125.4024
60. Von Ungern-Sternberg, A. (2022). Discriminatory AI and the law. In Cambridge University Press eBooks (pp. 252–278). https://doi.org/10.1017/9781009207898.020
61. Vredenburgh, K. (2021). The right to explanation. Journal of Political Philosophy, 30(2), 209–229. https://doi.org/10.1111/jopp.12262
62. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation. International Data Privacy Law, 7(2), 76–99. https://academic.oup.com/idpl/article/7/2/76/3860948
63. Wachter, S., Mittelstadt, B., & Russell, C. (2017). Counterfactual explanations without opening the black box: automated decisions and the GDPR. Harvard Journal of Law & Technology, 31, 841. https://doi.org/10.2139/ssrn.3063289
64. World Bank. (2019). Behavioral insights for tax compliance. World Bank. https://doi.org/10.1596/33046
65. Yismaw, A., Johns, A., Hussain, I., Creem, E., & Risler, M. (2021). Enhancing Return Risk Assessment for Examination: Recent DIF (Discriminant Function) Model Updates. IRS Research Bulletin. https://www.irs.gov/pub/irs-prior/p1500--2022.pdf
66. Zalcewicz, A. (2023). New technologies in the control of public finances and building public confidence in the state. Białostockie Studia Prawnicze, 28(2), 23–35. https://doi.org/10.15290/bsp.2023.28.02.02
67. Zheng, S., & Penetrante, M. O. (2025). Design and analysis of an AI-Driven Tax Avoidance Detection System in big data environments for public sector tax administration. International Journal of Research and Innovation in Social Science, IX(V), 3118–3130. https://doi.org/10.47772/ijriss.2025.905000241