

Auditing the Black Box: Rethinking Audit Assurance and Professional Liability in the Age of Artificial Intelligence

Dr. Gaduga Godwin, Esq

Lecturer at Accra Institute of Technology Adjunct Lecturer at Blue Crest University College

ABSTRACT

Artificial intelligence has moved from the periphery of audit practice to its operational core, and the profession's conceptual apparatus has not kept pace. This article examines what the opacity of machine learning systems does to two foundations of the auditing discipline: the assurance model, which rests on the auditor's ability to obtain and evaluate sufficient appropriate evidence, and the liability regime, which rests on a standard of care calibrated to human judgment. Drawing on auditing scholarship, standard-setting developments at the IAASB and the PCAOB, comparative regulatory instruments including the European Union's Artificial Intelligence Act, and the common law of auditor negligence, the article argues that neither wholesale prohibition of opaque systems nor uncritical reliance on them is defensible. It proposes a graduated algorithmic reliance framework that ties the permissible depth of reliance on an AI system to the demonstrable explainability of that system, the materiality of the assertion it supports, and the auditor's capacity to corroborate its output through independent means. The article also reformulates the negligence standard around the figure of the competent hybrid auditor and considers how liability should be allocated among audit firms, technology vendors, and audited entities. Particular attention is given to the position of developing economies, where regulatory capacity constraints sharpen every one of these questions.

KEYWORDS: Artificial Intelligence, Audit Assurance, Professional Liability, Explainability, Audit Evidence, Negligence, Forensic Accounting, Algorithmic Accountability.

ARTICLE DETAILS

Published On:
15 August 2026

Available on:
<https://ijiissh.com/>

1. INTRODUCTION

When the Public Company Accounting Oversight Board surveyed the largest audit firms in 2024, it found that generative artificial intelligence had already entered the audit workflow, initially in administrative and research tasks but with declared ambitions extending to risk assessment, scoping, and the evaluation of financial statement disclosures (PCAOB, 2024). The finding confirmed what practitioners had known informally for several years. The question facing the profession is no longer whether auditors will use artificial intelligence but whether the conceptual foundations of auditing can absorb a technology whose internal reasoning is, in many of its most powerful forms, inaccessible to the very professionals who deploy it.

That inaccessibility is what this article calls the black box problem, and it is not a rhetorical flourish. Burrell (2016) distinguished three forms of algorithmic opacity: opacity as corporate secrecy, opacity as technical illiteracy, and opacity intrinsic to the scale and structure of machine learning itself. The third form is the most troubling for auditing, because it cannot be cured by disclosure or by training. A deep neural network trained on millions of journal entries may flag a transaction as anomalous for reasons that no human, including its developers, can fully reconstruct. An auditor who relies on that flag, or on its absence, is relying on a process she cannot inspect in the way she can inspect a confirmation letter, a bank statement, or the arithmetic of a depreciation schedule. The difficulty runs in both directions. Auditors increasingly use opaque systems as tools of the audit, and they increasingly confront opaque systems as objects of the audit, embedded in the client's own estimation processes, credit models, and automated controls. In each posture, the profession's two load-bearing structures come under strain. The first is the assurance model itself, which since the codification of the international standards has rested on the requirement that the auditor obtain sufficient appropriate audit evidence as the basis for an opinion (IAASB, 2019). The second is the professional liability regime, which in most common law jurisdictions measures the auditor's conduct against the care that a reasonably competent member of the profession would have exercised in the circumstances. Both structures presuppose a human evaluator whose reasoning can be examined, documented, and, when necessary, litigated. Neither was designed for a professional who delegates portions of that reasoning to a statistical artefact.

This article makes three claims. First, the black box problem is best understood as an evidentiary problem rather than a technological one, and the existing architecture of the International Standards on Auditing already contains the conceptual resources to address it, provided those resources are extended deliberately rather than left to ad hoc interpretation. Second, the negligence standard in auditor liability requires reformulation, not abandonment. The reasonable auditor of the coming decade is neither the technophobe who refuses algorithmic assistance nor the enthusiast who accepts machine output without challenge, and the law should say so explicitly. Third, the allocation of responsibility among audit firms, technology vendors, and audited entities is currently indeterminate in ways that invite both under-deterrence and opportunistic litigation, and this indeterminacy falls hardest on jurisdictions, including many in Africa, whose regulators and courts have the least institutional capacity to resolve it case by case. The argument proceeds as follows. Section 2 traces the movement from conventional computer-assisted audit techniques to machine learning and generative systems, because the liability analysis depends on distinctions that this history makes visible. Section 3 develops the black box problem as a problem in the epistemology of audit evidence. Section 4 examines the response of standard setters and identifies the gap between the pace of technological adoption and the pace of normative guidance. Section 5 turns to professional liability and develops the standard of the competent hybrid auditor. Section 6 places the analysis in comparative regulatory context, with particular attention to developing economies. Section 7 sets out the graduated algorithmic reliance framework that constitutes the article's principal practical contribution. Section 8 draws out implications for forensic practice, professional education, and future research, and Section 9 concludes.

2. FROM COMPUTER-ASSISTED TECHNIQUES TO GENERATIVE SYSTEMS: THE CHANGING ANATOMY OF THE AUDIT

Auditors have used computers for half a century, and it would distort the analysis to treat every algorithm as a novel legal problem. The generalized audit software of the 1980s and 1990s, and the data analytics platforms that succeeded it, performed operations that were complicated but transparent. A recalculation routine, a three-way match, or a stratified sampling procedure can be specified in advance, tested against known inputs, and verified line by line. The auditor who relied on such a tool relied, in substance, on her own instructions executed at speed. Nothing in that arrangement disturbed the traditional account of audit evidence, and nothing in it disturbed the traditional account of professional responsibility, because the causal chain from auditor judgment to audit conclusion remained intact and inspectable.

Machine learning breaks that chain in a specific place. As Kokina and Davenport (2017) observed in one of the earliest systematic treatments of the subject, the distinguishing feature of contemporary artificial intelligence in auditing is that the system derives its decision rules from data rather than receiving them from the auditor. Issa, Sun, and Vasarhelyi (2016) mapped the audit tasks most susceptible to this form of automation and found that they clustered precisely in the areas the standards treat as judgment-intensive: risk identification, anomaly detection, and the evaluation of estimates. Appelbaum, Kogan, and Vasarhelyi (2017) reached a similar conclusion from the direction of research design, noting that the analytical techniques entering the audit had begun to outrun the evidential concepts available to describe them.

The practical trajectory has matched the scholarly prediction, although more slowly than the early literature anticipated. Salijeni, Samsonova-Taddei, and Turley (2019) documented how the large firms embedded data analytics in engagement methodology while struggling to articulate what those analytics contributed to the sufficiency of evidence. Krieger, Drews, and Velte (2021) found that adoption within firms was uneven and that the obstacles were organizational and epistemic rather than technical, with engagement teams uncertain how regulators and internal quality reviewers would treat conclusions resting on analytical output. The PCAOB's 2024 outreach confirmed both the acceleration and the caution: firms reported investment in generative tools alongside explicit anxiety about the auditability of the underlying data and the propensity of such tools to produce false or misleading content (PCAOB, 2024).

Generative systems deserve separate mention because they introduce a failure mode that classical machine learning does not share. A classification model may err, but it errs within the space of its defined outputs. A large language model can fabricate, producing fluent and confident text unmoored from any underlying record. For a profession whose entire epistemic warrant is the connection between assertion and evidence, fabrication is not merely an accuracy problem. It is a category threat. An audit memorandum drafted by a generative system may contain citations to workpapers that do not exist, and unless the reviewing auditor checks every reference, the documentation trail that liability regimes and inspection regimes both depend upon becomes unreliable at its source. The anatomy of the modern audit therefore contains three technologically distinct strata. Deterministic tools execute human instructions and raise no novel questions. Predictive machine learning systems generate conclusions through processes that are statistically validated but not transparently reasoned. Generative systems produce artefacts, including narrative artefacts, whose fidelity to the underlying record cannot be assumed. Any framework for assurance or for liability that fails to distinguish these strata will either overregulate the harmless or underregulate the hazardous, and much of the current commentary fails in exactly this way.

3. THE BLACK BOX PROBLEM AND THE EPISTEMOLOGY OF AUDIT EVIDENCE

Audit evidence is a defined term of art. Under the international framework, information becomes evidence when the auditor evaluates it and uses it to support the conclusions on which the opinion is based, and the auditor must judge both its sufficiency, a measure of quantity, and its appropriateness, a measure of relevance and reliability (IAASB, 2019). The definition presupposes an evaluating mind. It presupposes, further, that the reliability of information can be assessed by reference to its source and its nature, which is why evidence obtained directly by the auditor has traditionally ranked above evidence obtained from the entity, and documentary evidence above oral representation. The entire hierarchy is an epistemology in miniature, refined across a century of practice and litigation.

Where does the output of an opaque model sit within that hierarchy? The question resists a clean answer, and the resistance is instructive. The output is generated by the auditor's own tool, which suggests high reliability under the traditional source-based ranking. Yet the process generating it cannot be inspected, which under the traditional nature-based ranking suggests the opposite. A risk score produced by a neural network is, in evidentiary character, closer to the opinion of an expert whose reasoning is undisclosed than to the recalculation it superficially resembles. The standards have long required auditors who use the work of an expert to evaluate the expert's competence, objectivity, and methods. No parallel requirement yet attaches, with equivalent specificity, to the algorithmic expert that sits inside the firm's own software.

The problem deepens when the behavioral evidence is considered. Commerford, Dennis, Joe, and Ulla (2022) demonstrated experimentally that auditors discount contradictory evidence more readily when it comes from an artificial intelligence system than when it comes from a human specialist, a phenomenon the authors situated within the broader literature on algorithm aversion. The finding cuts against the popular assumption that the danger is blind deference to the machine. In estimation contexts, the greater danger may be unwarranted dismissal, with auditors overriding algorithmic warnings precisely when those warnings are most valuable. Brown-Liburd, Issa, and Lombardi (2015) had earlier identified the complementary risk: information overload and ambiguity in large data environments degrade rather than improve auditor judgment unless the decision environment is deliberately structured. Between aversion and overload, the human component of the human-machine dyad is demonstrably fallible in ways that the assurance model does not yet acknowledge.

It is tempting to answer the black box problem with a technological prescription, and a substantial computer science literature now advocates either post hoc explanation techniques or the outright substitution of interpretable models in high-stakes settings (Rudin, 2019). The prescription has force in auditing, but it cannot be adopted wholesale. Post hoc explanation methods produce approximations of a model's behavior, and an approximation of a reasoning process is not the reasoning process, a distinction that matters enormously when the explanation is later deployed in litigation as though it were a contemporaneous record of the basis for a professional judgment. Interpretable models, for their part, sometimes purchase transparency at the price of predictive power, and an auditor who chooses a weaker but explainable fraud detection model over a stronger but opaque one has made a trade-off that the standards currently give her no vocabulary to justify or even to record.

The epistemological conclusion is therefore modest but firm. Algorithmic output is neither inadmissible nor self-certifying. It is a species of evidence whose reliability depends on factors the traditional hierarchy does not capture: the provenance and representativeness of training data, the validation regime applied to the model, the stability of its performance across populations, and the availability of independent corroboration. Treating those factors as evidentiary attributes, assessable and documentable like any others, is the move that brings the black box inside the audit rather than leaving it as an unexamined input. Section 7 operationalizes this move. Before that, it is necessary to examine what the standard setters have and have not done.

4. RETHINKING ASSURANCE: SKEPTICISM, SUFFICIENCY, AND THE STANDARDS GAP

The international standard-setting response has been deliberate to the point of caution. ISA 315 as revised in 2019 modernized the risk identification standard and acknowledged automated tools and techniques, and the quality management standards that took effect in 2022 obliged firms to establish governance over the technological resources used in engagements (IAASB, 2019, 2020). The revision of ISA 500 on audit evidence, exposed for comment in 2023, proposed a reference framework intended to be adaptable to technology without being prescriptive about it, reflecting the board's judgment that principles outlast techniques (IAASB, 2023). At the time of writing, the revision remains in progress within a broader project on audit evidence and risk response, which means that the profession is conducting audits with generative systems under an evidence standard drafted before such systems existed in usable form.

The regulatory posture in the United States is similarly watchful. The PCAOB has placed technology on its research agenda and has stated that its existing standards are not viewed by firms as impediments to the use of generative artificial intelligence, while simultaneously emphasizing that due professional care and professional skepticism cannot be delegated to a tool (PCAOB, 2024). The United Kingdom's Financial Reporting Council reached a comparable position in its thematic review of audit technology, welcoming innovation while cautioning that firms had not consistently established how analytical tools contributed to the evidence

obtained (FRC, 2020). The pattern across jurisdictions is consistent. Regulators decline to prohibit, decline to prescribe, and rely on general principles whose application to opaque systems is left to firms, inspectors, and eventually courts.

There is wisdom in principles-based restraint, and this article does not argue for a technology-specific rulebook that would be obsolete before its ink dried. The difficulty is that the general principles, as currently articulated, do not answer the questions engagement teams actually face. Consider professional skepticism, the profession's most invoked and least specified virtue. Skepticism toward a management representation has a known grammar: corroborate, test the underlying records, consider incentive and opportunity. What is the grammar of skepticism toward a model? An auditor cannot cross-examine a gradient boosting machine. She can, however, interrogate its validation history, probe its performance on cases with known outcomes, examine whether the population on which it was trained resembles the population to which it is applied, and design procedures that would detect its characteristic failure modes. None of this is beyond the profession's competence, but none of it is currently required, described, or standardized, and quality review practices consequently vary in ways that inspection findings are beginning to expose.

The sufficiency concept requires similar renovation. Sufficiency has always been understood as a function of risk and of the quality of individual items of evidence, with lower quality demanding greater quantity. Algorithmic evidence complicates the calculus because its quality is not fixed. A model's reliability is a distribution, strong in the center of its training experience and weak at the edges, and the transactions that matter most in an audit, the unusual, the period-end, the management-initiated, live disproportionately at the edges. A sufficiency judgment that treats a model's aggregate accuracy as if it applied uniformly to every item the model touches commits a statistical error with legal consequences, because it is precisely the edge cases that generate restatements and lawsuits. The assurance model does not need new foundations to accommodate this insight, but it does need the insight stated authoritatively, so that the engagement partner who documents differential reliance across the risk spectrum is following guidance rather than improvising.

Munoko, Brown-Liburd, and Vasarhelyi (2020) argued that the ethical implications of intelligent systems in auditing extend beyond accuracy to accountability, and framed the question that standard setting has yet to answer: when an algorithm participates in a professional judgment, who answers for the judgment? The standards say, correctly, that the auditor does. The next section examines what that answer costs, and what it should cost, when the auditor's reliance was reasonable, and the machine was wrong.

5. PROFESSIONAL LIABILITY IN THE AGE OF ALGORITHMIC RELIANCE

5.1 The Negligence Standard and the Competent Hybrid Auditor

The common law of auditor negligence was built for human error, and its landmarks reflect that heritage. Lopes LJ's dictum in *In re Kingston Cotton Mill Co (No 2)* (1896) that the auditor is a watchdog and not a bloodhound fixed the standard at reasonable care and skill, not omniscience. *Ultramares Corporation v. Touche* (1931) confined the class of claimants in the American tradition out of concern for liability in an indeterminate amount for an indeterminate time to an indeterminate class, and *Caparo Industries plc v Dickman* (1990) performed a similar confinement in the English tradition through the requirements of proximity and purpose, building on the assumption of responsibility analysis in *Hedley Byrne & Co Ltd v Heller & Partners Ltd* (1964). What none of these authorities contemplated is a defendant whose alleged breach consists in trusting, or in failing to trust, a statistical system that no witness can fully explain.

The negligence inquiry asks how a reasonably competent auditor would have acted. Artificial intelligence destabilizes the inquiry at both ends. If algorithmic tools become ordinary professional equipment, then the auditor who declines to use them may herself fall below the standard, in the same way that a doctor who ignores available diagnostic imaging cannot plead fidelity to the stethoscope. Conversely, the auditor who uses such tools cannot discharge her duty merely by using them well in a technical sense, because the duty of care attaches to the professional conclusion, not to the operation of the instrument. The emerging standard, which this article names the competent hybrid auditor, has three components. She selects tools whose validation and limitations she understands at the level of an informed professional user, not a developer. She calibrates her reliance to the demonstrated reliability of the tool for the specific population and assertion at issue. And she preserves independent grounds for the conclusions that matter most, so that no material judgment rests on algorithmic output alone.

The behavioral literature gives this formulation empirical content that courts will eventually need. If auditors systematically discount algorithmic warnings in complex estimate settings (Commerford et al., 2022), then a plaintiff who shows that the firm's own model flagged the misstatement and that the engagement team overrode the flag without documented justification has a powerful case, and firms should assume that discovery will target model logs precisely because such logs now exist. Symmetrically, a team that accepted a model's clean output over contradictory traditional evidence will find that the model's aggregate accuracy statistics offer little shelter, because the standard of care is applied to the judgment made on the engagement, not to the tool's performance in general. In both postures, contemporaneous documentation of the reliance decision becomes the fulcrum of liability, which is an additional and independent reason why the assurance framework must specify what such documentation contains.

5.2 Allocating Responsibility Among Firm, Vendor, and Client

The auditor's non-delegable responsibility for the opinion does not settle the further question of how losses should ultimately be distributed when an algorithmic failure contributes to an audit failure. Three candidate bearers exist, and the current law reaches each of them imperfectly. The audit firm stands first in line, both doctrinally and practically, because it owes the professional duty and signs the report. The technology vendor stands behind contractual walls, since audit software is typically licensed under terms that disclaim fitness for any particular engagement and cap liability at the license fee, and the vendor owes no duty in negligence to the audited entity's shareholders under orthodox proximity analysis. The client stands in an ambiguous position, especially where the opaque system under audit is the client's own, embedded in its estimation processes or automated controls, and where the client's data practices shaped the model that misled the auditor.

This distribution is defensible as a matter of investor protection and indefensible as a matter of incentive design. Concentrating liability on the firm preserves the clarity of the assurance promise, but it also means that the party with the greatest capacity to prevent model defects, the vendor, bears the least exposure to their consequences. The European Union's product liability revision, which extends the product concept to software, signals one corrective direction, and the risk management obligations that the Artificial Intelligence Act imposes on providers of high-risk systems supply another, since compliance failures by a vendor will furnish evidence of defect and perhaps anchor contribution claims (European Union, 2024). Audit firms, for their part, should treat vendor contracts as risk allocation instruments rather than procurement formalities, negotiating audit rights over model documentation, warranties as to validation, and indemnities scaled to the engagements in which the tool will be used. Professional indemnity insurers are already asking these questions, and firms without answers will discover the cost of silence in their premiums before they discover it in a courtroom.

Where the opaque system belongs to the client, the analysis shifts from contribution to cooperation. An auditor cannot evaluate what she cannot access, and management that deploys machine learning in financial reporting processes assumes, on any sensible reading of its representations, an obligation to provide the documentation, validation records, and access necessary for the auditor to understand the system's role in the figures. Refusal or incapacity should be treated as a scope limitation with the reporting consequences that follow, and audit committees should understand that adopting inexplicable systems in the reporting chain is a governance decision with assurance consequences, not merely an operational one. The forensic dimension is immediate: in post-failure investigations, the first question is increasingly whether the numbers were produced by a process anyone in the organization could explain, and the answer allocates blame long before pleadings are drafted.

5.3 The Expectation Gap Revisited

The audit expectation gap, the durable divergence between what the public believes audits provide and what the standards actually require, predates artificial intelligence by decades, and Power (1997) diagnosed its structural roots in the ritual character of verification long before machine learning reached the engagement file. Artificial intelligence widens the gap from both sides. Public discourse presents the technology as superhuman in its detective capacity, inviting the inference that an AI-assisted audit should catch everything, and vendors marketing to the profession do little to discourage the inference. The realities of probabilistic detection, training data limitation, and adversarial manipulation ensure that AI-assisted audits will continue to miss cleverly concealed fraud, and each miss will now be accompanied by the question, asked in hindsight and in litigation, of why the machines did not find it. DeFond and Zhang (2014) showed that audit quality is already assessed by outsiders through coarse and outcome-based proxies. Algorithmic assistance will sharpen the coarseness, because outcome bias is easiest to indulge when the defendant had, or could have had, a tool.

The profession's response should be candor rather than mystique. Audit reports and firm communications should describe the role of technology in terms that neither conceal nor inflate it, and the standard setters should consider whether the auditor's report ought to disclose material reliance on automated tools in the same spirit that key audit matters disclose areas of significant judgment. Transparency of this kind does not eliminate the expectation gap, but it relocates the argument from what the jury imagines the technology could do to what the engagement actually documented it doing, which is the only ground on which the profession can win.

6. REGULATORY CONVERGENCE AND DIVERGENCE: THE VIEW FROM DEVELOPING ECONOMIES

The global regulatory environment into which audit AI is arriving is layered rather than unified. The European Union's Artificial Intelligence Act establishes a horizontal, risk-tiered regime whose obligations of transparency, human oversight, and technical documentation will reach some systems used in and around financial reporting (European Union, 2024). The United States has preferred sectoral guidance and voluntary architecture, most prominently the National Institute of Standards and Technology's risk management framework, which supplies a shared vocabulary of trustworthiness characteristics without binding force (NIST, 2023). UNESCO's recommendation on the ethics of artificial intelligence, adopted by consensus of its member states, articulates principles of transparency, accountability, and human oversight at the broadest level of generality (UNESCO, 2021). Jobin, Ienca, and Vayena

(2019) demonstrated that such instruments converge impressively on principles and diverge just as impressively on interpretation, enforcement, and allocation of responsibility, which is precisely the layer at which auditors and their insurers live.

For developing economies, the stakes are distinctive and, in this author's view, underexamined. Consider Ghana, whose audit profession operates under a modern companies statute, the Companies Act, 2019 (Act 992), and a data protection regime, the Data Protection Act, 2012 (Act 843), but without any instrument addressed to algorithmic systems in financial reporting or assurance. Ghanaian firms affiliated to global networks import engagement technology, and with-it engagement methodology, developed and validated elsewhere, on data generated in economies whose transaction patterns, informality levels, and documentary practices differ materially from local conditions. A fraud detection model trained principally on North American and European ledgers embeds assumptions about what normal looks like, and the edge-case fragility described in Section 4 is most acute exactly where local conditions diverge from the training distribution. The result is a quiet epistemic dependency: the reliability judgments that the standards require of the local engagement partner presuppose validation evidence that sits in another jurisdiction, in another company, behind another contract.

Regulatory capacity compounds the problem. Audit oversight bodies in much of Africa are young, thinly resourced, and already stretched by conventional inspection responsibilities, and it is unrealistic to expect them to develop independent model inspection capability in the near term. The realistic path is threefold. National regulators and professional institutes should adopt, by reference, international guidance on technology in the audit as it matures, rather than drafting from scratch. Regional cooperation, through bodies such as the Pan African Federation of Accountants, should pool the scarce technical expertise that no single oversight authority can retain alone. And local standard setters should insist, as a condition of network methodology deployment, on documented revalidation of imported models against local data, a requirement that costs the global firms little and purchases for local markets the one thing the technology cannot import, which is relevance. Scholars in the region have an accompanying duty to generate the empirical literature on which such requirements can stand, since the behavioral and archival research canvassed in this article draws overwhelmingly on data from mature markets.

7. A GRADUATED ALGORITHMIC RELIANCE FRAMEWORK FOR ASSURANCE ENGAGEMENTS

The analysis so far yields design requirements for a workable framework. It must distinguish the technological strata identified in Section 2, treat algorithmic reliability as an evidentiary attribute in the manner of Section 3, give skepticism and sufficiency the operational grammar demanded in Section 4, and generate the documentation on which the liability analysis of Section 5 turns. The graduated algorithmic reliance framework proposed here does this through a single organizing principle: the permissible depth of reliance on an algorithmic system in an assurance engagement is a joint function of the system's demonstrated explainability and validation, the materiality and risk of the assertion the system's output supports, and the availability of independent corroboration. The framework operates in five stages. The first is classification, in which the engagement team identifies every algorithmic system whose output will inform the audit, whether as a tool of the audit or object of it, and assigns each to one of the three strata: deterministic, predictive, or generative. Deterministic tools exit the framework at this stage, subject only to conventional testing of their operation. The second stage is interrogation, in which the team obtains and evaluates the system's validation dossier: training data provenance, performance metrics disaggregated by relevant subpopulation, known failure modes, version history, and the results of any independent testing. A system for which no adequate dossier can be obtained is treated, for evidentiary purposes, in the way an expert who refuses to disclose methods would be treated, which is to say its output may inform inquiry but cannot support a conclusion.

The third stage is calibration, in which reliance is set assertion by assertion. For low-risk assertions supported by well-validated predictive systems operating within their training distribution, output may be accepted subject to sample-based corroboration. For significant risks, algorithmic output may direct attention and shape procedures, but every material conclusion must additionally rest on evidence whose reliability is assessable by traditional means. Generative output receives a categorical rule at this stage: it may draft, summarize, and organize, but no representation of fact in the audit file may rest on generative production without human verification against the underlying record, and the verification must itself be documented. The fourth stage is challenge, the operational form of skepticism, in which the team performs at least one procedure per significant system designed to catch the system failing: seeded anomalies, back-testing against known outcomes from prior periods, or comparison of model output with an independent expectation developed without the model. The fifth stage is documentation, in which the reliance decision, the dossier evaluation, the calibration, and the challenge results are recorded in a form sufficient to permit an experienced auditor, having no connection with the engagement, to understand not only the conclusions but the division of epistemic labor between human and machine that produced them.

Three features of the framework deserve emphasis. It is standard-compatible, in that every stage can be expressed in the existing vocabulary of the ISAs and, with modest drafting, incorporated into the ISA 500 revision now in progress or into application material under the quality management standards. It is liability-relevant, in that a firm able to produce the documentation the fifth stage requires will have constructed, in advance, the very record that the competent hybrid auditor standard of Section 5 would examine,

converting a litigation vulnerability into a defense. And it is scalable, in that the intensity of stages two through four rises with risk and with opacity, so that smaller firms and smaller engagements are not crushed by governance built for global networks, a scalability that matters greatly for the developing economy contexts discussed in Section 6.

8. IMPLICATIONS FOR FORENSIC PRACTICE, EDUCATION, AND RESEARCH

For forensic practitioners, the black box problem is simultaneously a subject of investigation and a condition of work. Investigations of financial statement fraud will increasingly encounter records produced or filtered by algorithmic systems, and the investigator's first analytical task will be to reconstruct what the systems were shown, what they flagged, and what humans did with the flags. Model logs, override records, and validation files belong on standard document request lists now, and forensic reports should address algorithmic provenance with the same rigor long applied to documentary provenance. At the same time, forensic analysts deploying their own machine learning tools face, in the courtroom, an intensified version of everything this article has said about explainability, because expert evidence regimes in most jurisdictions require the expert's reasoning to be disclosed and testable, and an opinion that reduces to the model said so may not survive a competent admissibility challenge. Interpretability is therefore not a stylistic preference in forensic work but a condition of testimonial usefulness, which is an applied vindication of Rudin's (2019) argument for interpretable models in high-stakes settings.

For educators, the implication is curricular and cannot be met by appending a software module to an unchanged syllabus. The competent hybrid auditor described in Section 5 needs statistical literacy sufficient to read a validation dossier, conceptual understanding of how learning systems fail, and the professional judgment to integrate algorithmic and traditional evidence, and these capacities must be examined, not merely mentioned. Professional bodies in developing economies face the sharper version of this task, since their members will encounter imported technology before local guidance exists, and the qualification syllabus may be, for a period, the only instrument of preparation the jurisdiction possesses. There is also an obligation of scholarly self-awareness: the behavioral findings on which parts of this article rest await replication outside North American and European samples, and the region's universities are the natural site for that work.

The research agenda follows from the argument's open joints. The framework of Section 7 requires field testing, and its calibration stage in particular would benefit from archival study of how differential reliance decisions correlate with inspection findings and restatements. The liability analysis of Section 5 rests on doctrinal projection that only decided cases will confirm, and the first generation of AI-related audit litigation should be studied as it emerges with the attention the first wave of data breach litigation received. The comparative question of Section 6, how epistemic dependency in audit technology affects audit quality in importing jurisdictions, has to this author's knowledge no empirical literature at all, and it may be the most consequential of the three.

9. CONCLUSION

The auditing profession has survived every previous technology by absorbing it into an unchanged promise: an independent, competent human examined the evidence and formed an opinion. Artificial intelligence tests that promise more searchingly than its predecessors because it participates in the examining, not merely in the gathering, and because its participation is, in its most powerful forms, opaque to the participant it assists. This article has argued that the correct response is neither prohibition nor faith but structure. The black box can be audited, in both senses of that phrase, if algorithmic reliability is treated as an evidentiary attribute to be assessed, if skepticism toward systems is given the operational grammar it currently lacks, if the negligence standard is reformulated around the competent hybrid auditor who calibrates reliance rather than surrendering or refusing it, and if the resulting judgments are documented with a candor that will survive both inspection and cross-examination.

The graduated algorithmic reliance framework offered here is a contribution to that structure, deliberately built from materials the standards already contain so that its adoption requires evolution rather than revolution. What cannot wait for evolution is the question of responsibility. Courts, regulators, insurers, and audit committees are already allocating the costs of algorithmic failure by default, through contract boilerplate and doctrinal inertia, and default allocations are rarely just ones. The profession that claims the public's trust in the reliability of information now works alongside machines whose reliability it must itself establish. Meeting that obligation openly, with frameworks that can be inspected, taught, and litigated, is how auditing keeps its promise in the age of artificial intelligence, in Accra no less than in London or New York.

REFERENCES

- 1) Appelbaum, D., Kogan, A., & Vasarhelyi, M. A. (2017). Big data and analytics in the modern audit engagement: Research needs. *Auditing: A Journal of Practice & Theory*, 36(4), 1-27. <https://doi.org/10.2308/ajpt-51684>
- 2) Brown-Liburd, H., Issa, H., & Lombardi, D. (2015). Behavioral implications of Big Data's impact on audit judgment and decision making and future research directions. *Accounting Horizons*, 29(2), 451-468. <https://doi.org/10.2308/acch-51023>
- 3) Burrell, J. (2016). How the machine 'thinks': Understanding opacity in machine learning algorithms. *Big Data & Society*, 3(1), 1-12. <https://doi.org/10.1177/2053951715622512>

- 4) Caparo Industries plc v Dickman [1990] 2 AC 605 (HL).
- 5) Commerford, B. P., Dennis, S. A., Joe, J. R., & Ulla, J. W. (2022). Man versus machine: Complex estimates and auditor reliance on artificial intelligence. *Journal of Accounting Research*, 60(1), 171-216. <https://doi.org/10.1111/1475-679X.12407>
- 6) Companies Act, 2019 (Act 992) (Ghana).
- 7) Data Protection Act, 2012 (Act 843) (Ghana).
- 8) DeFond, M., & Zhang, J. (2014). A review of archival auditing research. *Journal of Accounting and Economics*, 58(2-3), 275-326. <https://doi.org/10.1016/j.jacceco.2014.09.002>
- 9) European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L, 2024/1689.
- 10) Financial Reporting Council. (2020). The use of technology in the audit of financial statements (Audit Quality Thematic Review). FRC.
- 11) Hedley Byrne & Co Ltd v Heller & Partners Ltd [1964] AC 465 (HL).
- 12) In re Kingston Cotton Mill Co (No 2) [1896] 2 Ch 279 (CA).
- 13) International Auditing and Assurance Standards Board. (2019). International Standard on Auditing 315 (Revised 2019): Identifying and assessing the risks of material misstatement. International Federation of Accountants.
- 14) International Auditing and Assurance Standards Board. (2020). International Standard on Quality Management 1: Quality management for firms that perform audits or reviews of financial statements, or other assurance or related services engagements. International Federation of Accountants.
- 15) International Auditing and Assurance Standards Board. (2023). Proposed International Standard on Auditing 500 (Revised): Audit evidence (Exposure draft). International Federation of Accountants.
- 16) Issa, H., Sun, T., & Vasarhelyi, M. A. (2016). Research ideas for artificial intelligence in auditing: The formalization of audit and workforce supplementation. *Journal of Emerging Technologies in Accounting*, 13(2), 1-20. <https://doi.org/10.2308/jeta-10511>
- 17) Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389-399. <https://doi.org/10.1038/s42256-019-0088-2>
- 18) Knechel, W. R., & Salterio, S. E. (2016). *Auditing: Assurance and risk* (4th ed.). Routledge.
- 19) Kokina, J., & Davenport, T. H. (2017). The emergence of artificial intelligence: How automation is changing auditing. *Journal of Emerging Technologies in Accounting*, 14(1), 115-122. <https://doi.org/10.2308/jeta-51730>
- 20) Krieger, F., Drews, P., & Velte, P. (2021). Explaining the (non-)adoption of advanced data analytics in auditing: A process theory. *International Journal of Accounting Information Systems*, 41, 100511. <https://doi.org/10.1016/j.accinf.2021.100511>
- 21) Munoko, I., Brown-Liburd, H. L., & Vasarhelyi, M. (2020). The ethical implications of using artificial intelligence in auditing. *Journal of Business Ethics*, 167(2), 209-234. <https://doi.org/10.1007/s10551-019-04407-1>
- 22) National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). U.S. Department of Commerce. <https://doi.org/10.6028/NIST.AI.100-1>
- 23) Power, M. (1997). *The audit society: Rituals of verification*. Oxford University Press.
- 24) Public Company Accounting Oversight Board. (2024). Spotlight: Staff update on outreach activities related to the integration of generative artificial intelligence in audits and financial reporting. <https://pcaobus.org/documents/generative-ai-spotlight.pdf>
- 25) Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206-215. <https://doi.org/10.1038/s42256-019-0048-x>
- 26) Salijeni, G., Samsonova-Taddei, A., & Turley, S. (2019). Big Data and changes in audit technology: Contemplating a research agenda. *Accounting and Business Research*, 49(1), 95-119. <https://doi.org/10.1080/00014788.2018.1459458>
- 27) *Ultramares Corporation v. Touche*, 174 N.E. 441 (N.Y. 1931).
- 28) UNESCO. (2021). Recommendation on the ethics of artificial intelligence. United Nations Educational, Scientific and Cultural Organization.