Tax Advisors in the Age of AI:
Professional Duty, Augmented Due Diligence and Cross-Border Liability

Joab García — Corporate Strategist

Abstract

The incorporation of artificial intelligence into international tax advisory practice raises a question that current doctrine has not adequately resolved: does delegating analytical functions to algorithmic systems dilute, preserve, or intensify the professional duty of the advisor? This article argues that it intensifies it — and does so in a structurally distinct way. Drawing on the ILA's due diligence doctrine, international accountability theory, and emerging AI governance frameworks, the article introduces augmented due diligence as a continuous professional governance function comprising four simultaneously operative dimensions: the non-delegable retention of accountability over algorithmic outputs; the proportional elevation of the due diligence standard; the exercise of genuine human oversight; and continuous compliance documentation. Failure to maintain any of these dimensions constitutes aggravated professional negligence whose reach extends across every jurisdiction the system touches.

I. The Advisor Behind the Algorithm

Artificial intelligence has become operational infrastructure in international tax practice. The OECD reports that AI adoption among tax administrations rose from 9 per cent in 2016 to 69 per cent in 2023.[1] Private advisory firms have followed — and in some respects led — that trajectory. Platforms now generate fiscal structures, model treaty exposures, score transfer pricing risk, and produce compliance memoranda faster than any human team. The efficiency gains are real and, in many contexts, the quality improvements are too.

What has not kept pace is the profession's understanding of what this shift means for professional liability. The instinctive response — that an advisor who reviews and approves an AI-generated opinion has discharged their duty — confuses the act of review with its substance. Signing off on a sophisticated algorithmic output is not the same as exercising independent professional judgment. It may, in fact, be its systematic substitution.

The risk has a name in the behavioural literature: *automation bias* — the tendency to defer to algorithmically generated recommendations without genuinely evaluating them.[2] In low-stakes contexts this produces inefficiency. In international tax advisory, where a single structural recommendation may affect clients across multiple

---

[1]OECD, Tax Administration 2025 (OECD Publishing, 2025). The report surveys 58 tax administrations. AI adoption rose from 9% in 2016 to 69% in 2023, with a further 24% reporting active implementation pipelines.

[2]S. Alon-Barkat and M. Busuioc, "Human–AI Interactions in Public Sector Decision Making: Automation Bias and Selective Adherence to Algorithmic Advice" (2023) 33 Journal of Public Administration Research and Theory 153. The study documents systematic deference to algorithmic outputs across professional contexts, including cases where the algorithmic recommendation was demonstrably incorrect.

jurisdictions, produce cascading compliance consequences, and in extreme cases constitute a predicate for money laundering offences, it can produce catastrophic professional exposure.

The argument of this article is straightforward. The use of AI in international tax advisory does not reduce the professional's duty of care. It redefines its content and raises its threshold, generating a governance obligation that existing doctrinal categories — due diligence, accountability, professional competence — each partially capture but none adequately describes in combination. To describe that obligation in integrated terms, this article proposes the concept of augmented due diligence.

II. Power and Pathology: AI in Tax Advisory

Three categories of AI tools now operate with practical significance in international tax advisory. Optimisation engines model multi-jurisdictional structures against treaty networks and domestic anti-avoidance rules to identify minimum effective tax positions. Predictive analytics platforms assign risk scores to proposed transactions based on enforcement data and comparative compliance records. Generative systems produce fully drafted advisory opinions, structural memoranda, and compliance documentation from advisor-defined parameters.

Each category offers genuine advantages. Speed, consistency, the capacity to survey dozens of treaty provisions simultaneously — these are not trivial gains. The problem is not the technology. It is the professional relationship to it.

Three structural risks deserve attention. The first is algorithmic opacity: sophisticated AI systems do not expose their reasoning in a form that a reviewing professional can audit. An advisor can verify that an output looks correct; they cannot always verify why the system produced it, or what it would produce under marginally different inputs. The second is error replication at scale: where a human advisor makes a mistaken assumption in one engagement, an AI system may reproduce that assumption across dozens of engagements simultaneously before anyone detects the defect. The third — and the one that carries the most direct legal consequence — is automation bias.

Research shows that deference to AI outputs is not merely a function of laziness or time pressure. It persists even when professionals are aware of the risk, even when they are instructed to maintain independent judgment, and even when the algorithmic recommendation is demonstrably wrong.[3] The advisor who approves an AI-generated structure without genuinely questioning its legal foundations does not merely risk a technical error. They risk establishing a pattern of conduct that, in a subsequent enforcement proceeding, will be characterised as a systematic abdication of professional responsibility — with full liability consequences attached to each instance of the pattern.

III. Augmented Due Diligence: A Governance Function, Not a Conduct Standard

Existing doctrine provides the building blocks but not the architecture. Due diligence, accountability, and professional compliance each address a dimension of what AI deployment requires from the tax advisor. None of them, taken alone, describes the full obligation. What is needed is a concept that integrates these dimensions into a single operative framework — one that functions continuously across the lifecycle of AI deployment rather than at discrete moments of professional judgment.

3.1 The Doctrinal Foundations

The ILA's Study Group on Due Diligence in International Law established that the content of the due diligence standard is not fixed — it adjusts to the nature of the risk, the capacity of the actor, and the potential scale of harm.[4] This contextual calibration is the key to understanding why AI deployment raises rather than preserves the professional standard. When an advisor uses a system capable of generating legally consequential recommendations across multiple jurisdictions simultaneously, the potential scale of harm is structurally different from anything the classical professional model contemplated. The standard adjusts accordingly.

The accountability dimension rests on a more fundamental principle: the obligation to answer for outcomes within one's sphere of control does not transfer with the delegation of a task.[5] The advisor who deploys an AI system to generate a planning structure retains full accountability for what that system produces — including errors, embedded biases, and legally deficient recommendations it generates autonomously. This is not a novel principle. It is the ordinary logic of vicarious responsibility applied to a new context. What is novel is the scale of what can go wrong within a single deployment decision.

The compliance dimension is grounded in the obligations that the FATF framework imposes on tax advisors as Designated Non-Financial Businesses and Professions.[6] Those obligations — client due diligence, transaction monitoring, suspicious activity reporting — do not disappear when they are executed by algorithmic systems. They follow the professional responsible for the process. Compliance, in this context, means

---

[4]T. Stephens and D. French, ILA Study Group on Due Diligence in International Law, Second Report (ILA, July 2016). The Report establishes that due diligence obligations are contextually calibrated: as the capacity to cause harm increases, so does the content of the standard required to satisfy the duty.
[5]Carolyn Moser, "The Concept of Accountability and Human Rights Violations" (2023) Max Planck Institute for Comparative Public Law & International Law Research Paper No. 2023-26. Moser distinguishes accountability from responsibility and liability, identifying its distinctive feature as the non-transferability of the obligation to answer for outcomes within one's sphere of control.
[6]FATF, Guidance for a Risk-Based Approach for the Accounting Profession (FATF, June 2019). Tax advisors qualify as Designated Non-Financial Businesses and Professions (DNFBPs) under FATF Recommendation 22 and are subject to client due diligence, transaction monitoring, and suspicious activity reporting obligations regardless of the tools through which those functions are performed.

governing the systems through which compliance is performed — not merely operating them.

3.2 The Concept

Augmented due diligence is the continuous professional governance function that arises when a tax advisor incorporates AI systems into their decision-making process. It integrates four simultaneously operative dimensions:

> First, the non-delegable retention of accountability over all outputs produced by systems under the advisor's control.
>
> Second, the proportional elevation of the due diligence standard, commensurate with the scale and jurisdictional reach of potential algorithmic harm.
>
> Third, the exercise of genuine human oversight — the capacity, grounded in professional expertise, to evaluate, question, and reject algorithmic outputs — as a condition of compliance.
>
> Fourth, the continuous documentation of that oversight process as a substantive obligation, not an administrative record.

Failure to exercise any of these dimensions constitutes aggravated professional negligence whose liability extends proportionally to every jurisdiction affected by the system deployed.

A reasonable objection: is this simply ordinary due diligence applied to a new tool? The ILA Study Group's own doctrine of contextual variability already accommodates evolving risk environments — one might argue that no new concept is needed. The difference, I would suggest, is structural rather than merely contextual. Classical due diligence is episodic: it asks whether the professional exercised adequate care at defined decision points. Augmented due diligence is continuous: it asks whether the professional maintained adequate governance across the entire lifecycle of a system capable of generating harm autonomously and at scale. That is not a difference of degree. It is a difference of kind.

3.3 Genuine Oversight and Its Counterfeit

The operationally most significant element of augmented due diligence is the distinction between *genuine* and *theatrical* oversight. Genuine oversight is oversight materially capable of influencing outcomes: it requires that the reviewing professional have enough domain expertise to detect when a system's output is legally deficient, factually inconsistent, or commercially implausible — and enough professional authority to reject or modify that output before it reaches the client.[7]

Theatrical oversight is its structural inverse. It is nominally present but incapable of detecting or correcting errors. The advisor who reviews an AI-generated opinion without the legal competence to question its assumptions, or who approves an output

---

[7]Kiteworks, "Human in the Loop: AI Compliance and Oversight Requirements" (March 2026). The analysis distinguishes between oversight "genuinely capable of influencing AI outcomes" and oversight that is "nominally present but structurally incapable of doing so" — a distinction the author terms theatrical oversight.

under time pressure without verification, has satisfied the form of professional review but not its substance. That distinction has direct liability implications: where oversight is theatrical, the act of review provides no defence against responsibility for errors the system produced.

The standard does not require the advisor to understand the algorithmic architecture of the system they deploy. That would be an impossible demand given the opacity of most commercial AI systems — and an unnecessary one. The standard is professional, not technical. What genuine oversight requires is intellectual independence sufficient to challenge a system's conclusions against the advisor's own expert knowledge of international tax law, treaty interpretation, and regulatory risk. Where that capacity is absent, the use of the system without independent validation constitutes negligence in itself, regardless of outcome.[8]

The EU AI Act embeds this principle in positive law. High-risk AI systems must be designed to enable effective human oversight, including information on the system's accuracy and the technical capacity for human intervention.[9] Tax advisory tools that generate legally consequential recommendations across multiple jurisdictions present a credible case for high-risk classification. Where that classification applies, the regulatory obligation reinforces the doctrinal one: oversight must be real, documented, and structurally capable of shaping outcomes. Anything less is compliance theatre.

IV. Liability Without Borders: Attribution in Cross-Jurisdictional AI Advisory

The jurisdictional complexity of AI-assisted tax advisory is best understood through a scenario that has moved from theoretical to routine. An advisor licensed in Jurisdiction A deploys an AI system hosted in Jurisdiction B to generate a planning structure for a client with operations in Jurisdictions C, D, and E. The structure is implemented. A revenue authority in Jurisdiction C later challenges it as aggressive tax planning and refers the matter for criminal investigation. The questions of applicable law, attribution of responsibility, and enforcement reach open simultaneously — and existing international private law frameworks were not designed to answer them.

4.1 The Multiplied Compliance Environment

The implementation of BEPS Pillar Two across more than forty jurisdictions has produced a compliance matrix of unprecedented density.[10] The GloBE rules require

---

[9]Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 on Artificial Intelligence (EU AI Act), Article 14. High-risk AI systems must enable effective human oversight, including the capacity for human intervention prior to or during deployment.

[10]OECD/G20 Inclusive Framework on BEPS, Global Anti-Base Erosion Model Rules (Pillar Two) (OECD, 2021, updated 2024–2026). As of 2025, over 40 jurisdictions have enacted or are actively implementing GloBE rules, creating a multi-layered compliance matrix that advisors increasingly navigate with AI-assisted tools.

advisors to navigate qualified domestic minimum top-up taxes, safe harbour provisions, and country-by-country reporting obligations across multiple regulatory environments at once. AI systems are deployed precisely to manage this complexity. The risk, less often acknowledged, is that errors in this environment do not stay local — they propagate across every jurisdiction where the structure operates.

When an AI-generated advisory produces regulatory exposure across multiple jurisdictions simultaneously, the applicable professional liability standard may differ in each. Each forum will apply its own definition of the duty of care, its own evidentiary rules for establishing breach, and its own enforcement mechanisms. The advisor faces a multiplicity of proceedings in which the content of their obligation is defined differently by each adjudicator.

4.2 Three Layers of Liability

The liability profile of AI-assisted tax advisory operates across three distinct legal layers, each with different enforcement mechanisms and different standards of proof.

At the **administrative and disciplinary level**, professional regulators may investigate and sanction advisors whose AI-assisted work fails to meet applicable competency standards. Professional bodies in multiple jurisdictions are progressively extending their supervisory mandates to encompass the governance of AI tools used in practice — a development driven less by formal rule changes than by the application of existing competency and independence standards to the new reality of algorithmic advisory work.

At the **civil level**, the advisor who deploys AI without genuine oversight may face liability for damages arising from deficient structures: tax penalties, interest charges, restructuring costs, and reputational harm to the client. The doctrinal question — whether the AI system is a cause of harm or the instrument through which the advisor's negligence operated — is being answered progressively in European jurisprudence and regulatory guidance in favour of the latter. The professional who controls a system that causes harm bears responsibility for what occurs under their supervision.

The **fiscal-criminal level** carries the gravest exposure. Advisory work that consistently generates structures later characterised as aggressive planning — executed through AI systems the advisor failed to supervise adequately — may support findings of recklessness or *dolus eventualis* in jurisdictions where those concepts apply. More critically: where deficient structures facilitate the concealment of taxable income or the mischaracterisation of revenue streams, the predicate conditions for money laundering offences may be met. The FATF framework places on tax advisors, as DNFBPs, an affirmative obligation to identify and report such risks.[11] An AI-generated output that appears compliant does not discharge that obligation.

4.3 A Framework Without a Forum

The Council of Europe's Framework Convention on Artificial Intelligence — the first binding international treaty on AI, opened for signature in September 2024 —

establishes baseline obligations of human oversight, transparency, and effective remedies across signatory jurisdictions.[12] The Convention does not directly regulate professional advisory services. But its articulation of human oversight as a structural requirement — not a best practice, not a recommended standard, but a binding obligation — reinforces the doctrinal foundation of augmented due diligence as an emerging cross-border standard of professional conduct.

What the Convention does not resolve — and what no existing instrument addresses — is the jurisdictional fragmentation of professional liability proceedings in cross-border AI advisory. A revenue authority that identifies a deficient structure can act within its own jurisdiction. It cannot directly regulate the foreign advisor whose system generated the structure. Mutual legal assistance mechanisms and joint audit programmes partially address evidentiary exchange; they do not coordinate professional liability proceedings across borders.

This gap is, for the moment, structural. What augmented due diligence offers in its absence is not a jurisdictional solution but a governance architecture that an advisor can apply consistently regardless of which forum eventually adjudicates their conduct. A documented record of genuine oversight — showing what the system produced, how the advisor evaluated it, what they questioned, and what they approved and why — is defensible before multiple regulatory audiences simultaneously. That is the practical value of the concept in the current enforcement environment.

V. What This Requires of the Profession

What follows has direct implications for the professional, regulatory, and academic communities — though not equally comfortable ones for each.

The first is that **AI does not create a new category of professional risk — it amplifies existing ones and changes their shape.** The duties of competence, independence, and client protection that have always governed tax advisory do not dissolve in algorithmic environments. They migrate, with intensified force, into the governance of the systems through which advisory functions are now performed. An advisor is not absolved by the sophistication of the tool. They are measured, in part, by the quality of the governance they exercise over it.

The second is that **the distinction between genuine and theatrical oversight will become the central evidentiary question in AI-related professional liability proceedings.** Courts and regulators will not ask whether a human was present in the advisory process. They will ask whether that human was capable of intervening meaningfully. Professional education, competency standards, and disciplinary frameworks have not yet incorporated this distinction with the clarity it requires. The

[12] Council of Europe, Framework Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law (opened for signature 5 September 2024). The Convention is the first binding international treaty on artificial intelligence, establishing baseline obligations of human oversight, transparency, and effective remedies applicable across signatory jurisdictions.

gap between the standard that augmented due diligence demands and the standard most practitioners currently apply is not marginal. It is structural.

The third is that **the advisor who exercises genuine oversight is not only better protected legally — they are providing a qualitatively different, and superior, professional service.** The capacity to critically evaluate AI-generated output, to identify its limitations, and to take documented responsibility for the final recommendation is not merely a risk management exercise. It is the professional's core contribution in an environment where the analytical work is increasingly algorithmic. The alternative — treating AI output as a finished product rather than a starting point for expert judgment — is not just a liability risk. It is a professional abdication.

The tax advisor of the next decade will not be distinguished by their ability to operate AI systems. Any firm with a budget can operate them. They will be distinguished by their capacity to govern them — to maintain the independence, the documented oversight, and the professional judgment that converts algorithmic output into trusted advice. That capacity is what augmented due diligence describes. The advisors who build it first will not merely reduce their liability exposure. They will define what professional competence means in an era when the analysis is increasingly algorithmic and the judgment — irreducibly — remains human.

Joab García

Corporate Strategist