

THE SCEPTIC'S GUIDE TO THE AI AUDIT APOCALYPSE: WHY YOU'RE RIGHT TO QUESTION (AND WHY YOU MUST ADAPT ANYWAY)

AUTHOR: AARON GREENMAN

If you are an internal auditor rolling your eyes at yet another article about Artificial Intelligence, you are not alone. There is a palpable fatigue in our profession, a feeling that we are living through a cycle of breathless hype that ignores the messy reality of our day-to-day work.

We have been here before. In the 1990s and 2000s, we were told that data analytics would revolutionise everything. Yet, decades later, many businesses and audit functions are still struggling to get clean data out of their ERP systems. The scepticism is valid. In fact, it is necessary. But while the hype might be overblown, the hazard of ignoring this shift is fatal. Here is the honest truth about AI in internal audit: why it's a headache today, but why it is the only way to survive tomorrow.

The Sceptic's Case: The Garbage In Reality

Let's recognise the elephant in the room: AI is often a solution looking for a problem.

- 1. The Garbage-In, Garbage-Out Wall:** The glossy brochures for AI tools including those developed for audit and assurance, assume organisations have pristine, structured data and well-defined processes. Most don't. As one seasoned practitioner noted, garbage-in, garbage-out is the iron law of AI. If you feed an AI model flawed, biased, or incomplete data, it doesn't just give you a wrong answer, it gives you a wrong answer with supreme confidence. Currently, progress in realising AI benefits is stalling in many areas simply because of poor data quality. Similarly, many teams lack well defined processes and measures making it difficult to define what good looks like.
- 2. The Black Box and the Death of Assurance:** How do you provide independent assurance on a decision you can't explain? This is the Black-Box problem. Deep learning models often cannot explain how they reached a conclusion. If an AI flags a vendor as high-risk, but cannot tell you why, can you professionally stand behind that finding? Furthermore, if internal audit uses the same (or alternate) AI models to assess risks and controls, we risk losing our independence.
- 3. The Hallucination Hazard:** Generative AI doesn't know facts. It predicts, it hallucinates and it regularly makes mistakes. This can create more work, not less, because every output requires a human-in-the-loop to verify it. If I have to fact-check every output, where is the efficiency gain?

4. The Water Cooler Effect: The most effective audit work often happens in the grey areas, the uncomfortable silences in an interview, the body language of a CFO, or the casual watercooler conversation six months before a project fails. AI cannot read the room. It lacks business context, ethical nuance, and emotional intelligence. There is a genuine risk that by outsourcing testing to AI, we weaken the auditor's judgment. We risk raising a generation of auditors who blindly trust the machine rather than developing the gut check that detects fraud.

The Pivot: Why the Burning Platform is Real


Despite these valid frustrations, the ground is shifting beneath us. The scepticism that protects us today could bury us tomorrow.

- 1. The Endangered Species List:** The World Economic Forum's Future of Jobs Report 2025 paints a brutal picture. Accountants and Auditors are among the top occupations expected to decline by 2030, sitting alongside data entry clerks. Conversely, roles like Strategic Advisors and Risk Management Specialists are projected to grow. The market is sending a clear signal: the Purist auditor, who focuses solely on retrospective assurance and compliance, is becoming obsolete. To survive, the profession must shift from checking-the-box to becoming a truly trusted advisor.
- 2. From Librarian to Journalist:** The nature of the technology is changing faster than our scepticism. Generative AI is like a librarian. You asked a question, and it fetches a book. The new Agentic AI however is more like an investigative journalist. It doesn't just retrieve; it seeks leads, connects dots across silos, and drafts reports. This kills the sampling defence. In a world where AI can analyse 100% of a population (up to billions of transactions) in near real-time, relying on a traditional sample of 25 is no longer just old-fashioned; it's negligent. As one expert put it: "AI removes the haystack and leaves us to focus on all the needles".
- 3. The Old Boring is the New Strategy:** The best use cases for AI aren't the sci-fi scenarios, they are the drudgery. AI excels at the tasks auditors hate: writing process narratives, harmonising risk and control matrices, and summarising regulatory changes. By outsourcing the boring work to the machine, the auditor is freed up to do the one thing the AI cannot: apply judgment, empathy, and strategic context. The goal isn't to replace the auditor, but to elevate them from a documenter to a critical thinker.

The Verdict: Be the Governor, Not the User

We are at a crossroads and witnessing the democratisation of assurance. If internal audit cannot provide insights deeper than what a manager can get from a chatbot, the function has limited value. We can continue to point out the flaws in the data and the risks of the black-box, or we can step into the vacuum and help leaders manage those very risks.

The hot topic for the future isn't just using AI; it is governing agentic AI systems. Organisations are rushing to adopt these tools with little oversight. Who is better positioned than Internal Audit to ask: Is this data biased? Is this model explainable? Is this secure? Is this assurance truly effective?



We have a unique opportunity to stop being the function that lags behind the business and start being the innovators. We can be the ones who tell the C-suite, "Yes, trust the system, but only once we have verified it".

Therefore, the auditor's role must shift from checker to validator, ensuring that the algorithms running the business aren't hallucinating, leaking proprietary data, or making biased decisions. We are no longer auditing the transaction; we are auditing the brain that processed the transaction.

The IIA's Vision 2035 report[1] also supports this radical shift: moving the profession from spending 76% of its time on assurance to a near-even split with advisory work (41% advisory). This requires Assurance-in-the-Loop, being part of the design phase of systems rather than grading the homework after the fact.

You don't have to love AI. You don't even have to trust it yet. But you cannot afford to ignore it. The train is leaving the station, with or without us. The choice is simple: evolve into a tech-enabled strategic advisor or prepare to be automated.

[1] <https://www.theiia.org/globalassets/site/foundation/latest-research-and-products/vision-2035-report.pdf>

**CONSULTANCY THAT
VALUES BIG THINKERS,
WHO SOLVE UNIQUE
PROBLEMS.**