

The Q3 Number Nobody Verified

Why your finance team's next forecast may already have been drafted by an AI nobody on your team can actually audit and what your role as Chief People Officer has to do with it.

By Christian E.N. Essoh | Atosenography

It is a Tuesday in November. Your CFO drops the Q3 board pack into the shared folder at 09:47, fourteen minutes before the meeting. You skim the executive summary, note that the forecast holds, and move on. The board accepts it. The chair signs off the quarterly update to investors that afternoon.

Three weeks later, a junior analyst on the FP&A team mentions, in passing, that she had used a generative AI tool to summarise the regional revenue inputs just to save time. The forecast itself was hers. But the framing, the narrative, and one of the assumptions had come from the model's output. She had not flagged it. Nobody had asked her to.

By the time her manager mentions it to your CFO, the quarterly update has been read by 340 institutional investors. Nobody knows what the model actually said. Nobody kept the prompt. Nobody verified the one assumption that anchored the rest. It is, almost certainly, fine. And there is no way for anyone in your organization to confirm that.

This Is Not a Hypothetical

Some version of this scene is already happening inside organizations of every size, in every sector. The specific function changes as in; finance, legal drafting, hiring shortlists, customer correspondence, strategic memos; but, the pattern is identical. A capable, trustworthy employee uses an AI tool to do part of their job. The output passes through their judgement, briefly, and into the work product of the organization. Nobody else sees the seam.

This is the practical face of an enormous, quiet shift that has happened to your workforce in the last twenty-four months: not the visible adoption of AI, which is well documented, but the invisible *integration* of AI into the cognitive substrate of how decisions get made. The question is no longer whether your people are using AI. They are. The question is whether they are maintaining control of the process or whether the process is quietly running them.

And the answer to that question is not a training problem. It is a measurement problem; and almost no organization has begun to measure it.

Why This Lands on Your Desk

If you are a Chief People Officer, the CPO role that emerged in the mid-2010s in technology firms and now sits in the C-suite of most serious enterprises, this might appear, at first, to be someone else's problem. The CFO owns the forecast. The Chief AI Officer, if you have one, owns AI governance. Compliance owns the audit. You own people: their capability, their development, their retention, their engagement, their performance.

Which is precisely why this lands on your desk. Because what is being compromised, slowly and invisibly, is not a system or a process. It is the capability of your people to do the work they were hired to do.

When a finance analyst routinely accepts the framing of a model without verifying it, she is not learning to do finance better. She is learning to do something else entirely; a hybrid task in which the part she actually owns is shrinking, week by week. Six months later, the work product still looks the same. Two years later, the person doing it no longer has the cognitive infrastructure to verify the model independently, even if she wanted to. She has been quietly trained out of the capacity that made her valuable in the first place.

Now multiply that by every analyst, every manager, every director in your organisation using AI tools every day. That is the workforce capability question. Not whether your people are using AI; but whether the cognitive habits that distinguish good judgement from compliance with model output are still there, and still developing, in the people you are paid to develop and retain.

This is your problem because nobody else in the organization is asking it.

The Evidence That Something Has Shifted

Three findings from the last eighteen months tell a coherent story.

First, AI adoption is uneven across your workforce in ways that intuition does not predict. In the United Kingdom, 80% of men in the workplace report using AI tools in their work. Among women, the figure is 59%. At director level, 73% of men received structured AI training from external providers compared with 58% of women in the same role.

Globally, a 2024 Harvard Business School meta-analysis by Otis, Delecourt, Cranney and Koning; synthesising eighteen studies covering approximately 143,000 individuals worldwide—found that women had 25% lower odds of using generative AI than men. The gap held across regions, sectors, and occupations. It persisted across income levels and education. It did not disappear even when access to the technology was equalized.

The standard interpretation of this data is that there is a gender gap in AI adoption, which there is. But the interpretation that matters to you, as CPO, is different and more interesting.

The same research shows, repeatedly, that the people adopting AI more slowly are not disengaged. They are the ones reporting the highest ethical reservations, the deepest concerns about model opacity, and the strongest instinct to verify before they trust. They are not behind. They are measuring what they see, and they are sceptical of it.

Read it that way and the picture inverts. The slower adopters are preserving exactly the cognitive habits; verification, justification, skepticism, that prevent AI from running an organization without its knowledge. The faster adopters, in many cases, are losing those habits without realising it.

The people pushing back the hardest may be the people
doing the most important work.

— The CPO question, restated

Why More Training Will Not Solve This

The conventional response, when a problem like this arrives on a CPO's desk, is more training. More AI literacy modules. More workshops. More certification. There is a multi-billion-pound industry ready to provide it, and most of it is competent.

It will not solve this problem.

Training tells people how to use the tools. It does not tell you whether the people you have trained are maintaining their independent judgement while they use them. Those are two different things. The first is a curriculum problem and is largely solved. The second is a measurement problem and is largely untouched.

If you train your finance team to use AI more effectively, you will almost certainly make the Q3 forecast scenario above *more* likely, not less. The analyst will use AI more confidently, more often, more deeply. The seam will become harder to see, not easier. Her capacity to verify independently will atrophy in proportion to how much she relies on the tool. That is the predictable consequence of training-without-measurement.

What you need, before or alongside any training programme, is an instrument that tells you whether the people in your organisation are maintaining the cognitive habits that make their judgement worth trusting. You need to be able to answer, for any given individual or team: are they working *with* AI, or are they being worked *by* it?

There has, until recently, been no instrument that could answer that question.

Atosenography — A Method, and an Instrument

Atosenography is the method for measuring how well your people maintain agency as AI integrates into their work and for closing the specific gaps the measurement reveals. It exists because the gap between what training delivers and what organizations need has become commercially significant, and is widening.

The intellectual foundation of the method is a published volume *UBIO EXIT*; which sets out the framework for understanding what is preserved, and what is lost, when human judgement integrates with systems designed to anticipate and assist it. The framework predates the current AI moment by some years. The current AI moment is what made it commercially urgent.

From that framework comes a software platform, **XeXina**, that performs the measurement at the scale an enterprise actually needs. It maps each person across twenty-eight cognitive dimensions, through three integrated systems. The output is a cognitive map, unique to the individual, that tells you and them where the capacity to maintain agency in AI-integrated work is strong, where it is weak, and where intervention will produce the largest gain.

Crucially, and this is what makes the conversation different from one about training; the recommendations that follow are built *from the measurement*. Not from a generic module. Not from a curriculum designed for a notional workforce. From this person's data, this team's patterns, this organization's specific shape.

The output is not abstract. It is operational. A team lead reads it. A CPO uses it. A regulator, increasingly, will ask for it.

The Regulator Will Get There Before You Do

Article 14 of the EU AI Act, which entered force in August 2024, mandates *meaningful human oversight* of high-risk AI systems. The Act is explicit that policy documents and training certificates do not satisfy the requirement. What is required is evidence that is auditable; individual-level evidence showing that the people exercising oversight are actually capable of doing so.

The core enforcement date for high-risk AI obligations, including Article 14, is 2 August 2026. By political agreement reached on 7 May 2026, a subset of Annex III systems, including those used in employment decisions, was extended to 2 December 2027. Either way, the deadline is now visible on every Chief AI Officer's calendar and on every CPO's, because the employment subset is the slice that lands closest to your function.

Most UK organizations are not preparing for this. The few that are preparing are doing so by writing more policy documents, which will not satisfy the requirement.

If your finance team produces a Q3 forecast based on an AI tool whose framing the analyst could not independently verify, and that forecast shapes a regulated disclosure or a high-stakes decision, you do not have a defensible record that human oversight took place. You have a

record that a human typed the final document.

Those are not the same thing. The regulator knows the difference. Increasingly, your customers know the difference. Your competitors who have begun measuring already have a head start on the evidence trail.

What This Looks Like, Operationally

A complete cognitive cartography for one individual takes sixty minutes. It is not a personality test, not a self-report questionnaire, not a productivity assessment. It is a measurement, conducted through behavioral telemetry that is capturing how the person actually thinks, not what they say they think.

From that single hour comes a twenty-eight-dimension cognitive report, delivered immediately, with industry-specific interpretation and targeted recommendations. For organizations, individual reports aggregate to team-level pattern analytics. For regulated organizations, the audit trail is produced in a format that maps directly onto Article 14 oversight evidence requirements.

Critically, and this is non-negotiable, written into the professional code that governs the discipline, is that the measurement is not designed for performance management. It cannot be used for firing or hiring decisions. It does not surveil. It does not score productivity. Every individual retains permanent ownership of their cognitive map. The instrument that allows you to evidence Article 14 compliance is the same instrument that gives your people sovereignty over the data being produced about them. That is built into the architecture, not bolted on.

If you are a CPO, this is the unusual combination you need: an instrument that produces enterprise-grade evidence without compromising the employee-first relationship that allows your function to work in the first place.

What to Do, This Quarter

XeXina is in controlled rollout. Access is by waitlist. The first twenty-five organisations to onboard receive Founding Customer status: preferential terms locked for twelve months, priority onboarding, direct access to the founder, and structured input on the platform's feature roadmap. Founding status is capped, time-limited, and assessed in order of application.

This is not a request to commit to an enterprise contract. It is an invitation to begin measuring; one team, one function, one pilot deployment before the regulator arrives, before the next Q3 forecast lands on the board pack, and before the people in your organization lose any more of the cognitive capacity you are paid to develop and retain.

The reading you have just completed takes most CPOs about twelve minutes. The decision to join the waitlist takes about ninety seconds. The cost of not doing so, measured in workforce capability decay over twenty-four months, is considerably larger than the cost of a Founding Customer engagement.

Join the XeXina Founding Customer Programme. The waitlist is open at atosecond.com. Founding places are limited to twenty-five organisations.

Atosenography is a trading identity of Easberg Ltd, registered in England and Wales (Company No. 14898760, D-U-N-S 230 676 246). XeXina is a B2B SaaS cognitive measurement platform. ATOSEN (ATS) is a utility token used to access platform services. The intellectual foundation of the method is set out in UBIO EXIT by Christian E.N. Essoh. UK adoption data referenced is published by HM Government. The global meta-analysis referenced is Otis, N. G., Delecourt, S., Cranney, K., & Koning, R. (2024). Global Evidence on Gender Gaps and Generative AI. Harvard Business School Working Paper No. 25-023 (revised August 2025). EU AI Act references draw on the official text and the political agreement of 7 May 2026 concerning the implementation timeline for Annex III high-risk systems. Nothing in this article constitutes financial advice.