

AI Adoption and the Cost We Aren't Measuring

AI adoption, workforce capability and the judgement gap in financial services. Building on the Financial Services Skills Commission's research for HM Treasury, and offered to the wider industry discussion.

The question the numbers don't answer

Ask a financial services firm how AI adoption is going and you will usually be given a number. Seats deployed, licences issued, use cases live and cost savings booked. It is a confident answer, but it does not tell you whether the organisation is becoming more capable because of it.

Those measures were not wrong. While the question was whether people were using AI they were the right things to count. They are the wrong measures if they are now read as evidence that the organisation is becoming more capable. They tell you how many people can open the tool, and nothing about whether anyone has become better at the work they are accountable for.

The Financial Services Skills Commission has made the point more precisely than the sector has absorbed it. In *A Workforce Transformed* it concludes that access to the technology itself is universal and not a differentiator.¹ If that is right, and we believe it is, the sector has built its adoption narrative around the one variable that will not distinguish anybody from anybody.

The same report sets out five challenges.² Three turn on a question almost no firm can answer: upskilling the workforce, redesigning routes in for early careers and closing the time to competence gap all require a firm to know whether its people are becoming more capable.

The Commission has already given firms a framework for identifying future skills and proficiency, but it is largely role agnostic. In an AI-enabled organisation, the capability required should reflect an individual's role, accountability and the judgement they are expected to exercise.

450,000

highly skilled people the sector must recruit and train by 2035 to replace those leaving or retiring, from a highly skilled workforce of 780,000³

FSSC, *A WORKFORCE TRANSFORMED*,
MAY 2026

30–50%

of tasks in most UK financial services roles expected to change through automation⁴

FSSC, *A WORKFORCE TRANSFORMED*,
MAY 2026

34%

of UK financial services firms report complete understanding of the AI they use⁵

BANK OF ENGLAND AND FCA AI SURVEY,
NOV 2024

The sector may need to recruit and train 450,000 highly skilled people over the next decade to replace those leaving or retiring, while AI is changing many of the tasks through which people historically became highly skilled.

The Commission has done the skills diagnosis, and done it well. This paper asks the question behind it. A skills gap concerns what people know and can often be addressed through training. A judgement gap concerns how experience and independent view are formed in the first place, and training alone does not close it.

Can we evidence that our people are becoming more capable as AI becomes more capable, and that the capability they develop is sufficient for the decisions, controls and accountability they carry? If AI performs more of the work through which professional judgement was traditionally developed, how will firms build the experience and judgement the next generation will need?

There is also a precedent worth remembering. Previous financial crises exposed what can happen when accountability outpaces the capability to understand and challenge the financial instruments, models and risk processes on which firms rely. AI risks creating a similar governance gap, but one that extends far beyond the Board.

We make four arguments.

- 1 Capability, not access, is the thing to measure, and almost no firm can evidence it.** Firms need to know whether people are becoming more capable, where capability is developing unevenly and whether the work is becoming better, not simply faster.
- 2 AI is creating an Accountability–Capability Gap.** The level of AI awareness and capability required is different for a Board member, accountable executive, control owner and practitioner. It should be proportionate to the accountability held, the consequence of the decision and the degree of reliance placed on AI. Senior leaders are being asked to govern and take accountability for a technology they may have limited understanding and practical experience of using.
- 3 AI can create capability debt by removing some of the experiences through which professional judgement develops.** The productivity benefit is immediate, while the cost may surface later through longer time to competence, weaker independent challenge and a reduced pipeline of experienced judgement.

4

The destination is human elevation, not human preservation. The question is not how much work AI should do. It will do more, and then more again. The question is what humans must continue to learn, know, decide and own, and how firms deliberately develop and assure that capability.

WHAT FIRMS SHOULD DO

We propose four actions. Define accountable AI capability by role from Board level to practitioner and test whether it exists. Change what you measure, from access and activity towards capability, quality and distribution. Give AI the institutional and individual context that makes its output useful and ownable. And deliberately rebuild the learning loop as AI removes some of the work through which judgement was previously formed, assuring the human, the AI and the process together.

WHY WE ARE WRITING THIS, AND HOW TO READ IT

Plenitude is a Supporter of the Financial Services Skills Compact and an early adopter of AI in our own firm. We are not writing as observers. We have put our methodology into our systems, measured adoption in ways that told us too little about capability, and changed our approach as a result. Section 3 is that account, including the parts that did not work.

We are strongly pro-adoption, and this is not an argument for AI doing less. Nor are we trying to predict every skill an AI-enabled workforce will need. We are asking whether firms are building sufficient AI awareness across their workforce and, more importantly, developing the human capability required as AI becomes more capable and embedded in work — to use it effectively, provide context, challenge its outputs, exercise professional judgement and discharge the accountability people retain. Passages marked our view are our judgement, and one claim in Section 2 is labelled a hypothesis because that is what it is.

The wrong scoreboard

Most firms we speak to can tell us how many people have access to an AI tool. Very few can tell us whether the work has improved.

Readers will have seen the announcements: ChatGPT, Claude or other tools rolled out to tens of thousands of employees. Deployment has become a headline; capability has not yet become the measure.

This is not a criticism of the people producing those numbers. Licence counts are easy to track and rise reliably, and they were among the first metrics we tracked ourselves. The problem is that nobody has since asked what they do not tell us: who is not using the tools, or whether the people carrying accountability for AI-influenced decisions could challenge one if they had to.

Financial services depends on decisions being explainable and accountable. AI makes both harder when the provenance of an AI-supported output cannot be reconstructed, or when the person accountable for the outcome lacks the capability to interrogate how AI influenced it.

What the sector actually measures

The evidence is not confined to the firms we speak to. The Bank of England and FCA's third survey of AI in UK financial services measures adoption through the proportion of firms using AI and the number of use cases in operation.⁵ Useful indicators of activity, but they do not tell us whether the people using the systems have become more capable. The closest it comes is asking firms how well they understand their own models, and even there only around a third report full understanding.

Time saved and return on investment matter, but they answer a different question. A firm can save thousands of hours without knowing whether its people are becoming more capable, less capable or simply more dependent on the technology. AI has dramatically reduced the cost of producing something that looks like work. It has not reduced the cost of deciding whether that work is any good.

The growing concern about "AI slop" is a good example. AI makes it easier to produce more, but more does not necessarily mean better. As the volume of AI-generated work increases, the ability to tell what is useful, what is wrong and what is worth acting on becomes more important, not less.

The Commission's own research captures the tension. One executive was refreshingly direct: return on investment was not yet the priority, the objective was to get people using the tools.⁶ Entirely rational during rollout, but rollout metrics should not become permanent measures of capability.

The measurement problem is about to become visible

Two separate pieces of work are easily confused. The Financial Services Skills Compact is a set of employer commitments; it launched in 2026 and is moving into its ongoing cycle.⁷ The HM Treasury research programme is separate: A Workforce Transformed was Phase One and assessed how technology is changing skills and the workforce. The remaining phases consider how the challenge should be addressed and who must act, and report in January or February 2027.⁸

One of the Compact's four commitments changes the measurement problem: signatories must publish annual progress. More than twenty regulated firms have signed, and each will shortly have to explain what progress means. There is no common definition of AI capability, or of what “upskilled” means in practice, so our expectation is that many will default to the numbers already available, which are not comparable. This is not a criticism of the Compact. A commitment to publish progress simply creates an immediate need for a definition of it.

There is a parallel worth borrowing from the wider skills debate. Falling training spend is often read as firms investing less in their people. It may equally mean a firm understands its gaps better and targets development more precisely. Training spend does not prove capability any more than licence deployment proves adoption maturity. Both measure an input, because the input is easy to count.

The measure the Commission already needs

Three of the Commission's five challenges depend on this, and one names the measure. Closing the time to competence gap requires measuring time to competence, which the sector does not do at scale.² It is where our argument and the Commission's agenda meet exactly, and it is the number we would expect AI to be moving quietly in the wrong direction while every visible adoption metric suggests the opposite.

OUR VIEW

If AI is removing some of the work through which people became competent, time to competence might be lengthening. We do not know whether it is. Very few firms measure it, and fewer measure it consistently enough to see a trend. It is a testable question, and it may be the most useful thing the sector could start measuring.

What the default costs

Many firms are clear that human judgement and expertise will remain important as AI takes on more of the work. Far fewer have considered how that judgement and expertise will continue to develop as the work through which they were traditionally built begins to change.

Adoption has typically been run as a technology programme and measured as one, while rapid advances in AI capability are outpacing organisational understanding and awareness. Four consequences follow.

ONE

Accountability may be outpacing AI awareness and capability

Accountability for a regulated activity rests with named individuals, and it does not transfer because AI performed part of the work. The Mills Review is explicit that as autonomy increases the human role shifts from operator through collaborator, consultant and approver to observer, and that accountability does not move with it.⁹ The AI Adoption Plan asks for AI capability across boards, specialists and frontline teams, naming boards specifically.¹⁰

The financial crisis precedent outlined earlier is relevant here. Boards have previously found themselves accountable for risk decisions increasingly dependent on instruments and models they could not themselves interrogate. AI risks creating a similar gap between accountability and the capability to exercise it, but one that extends beyond the Board into the decisions, controls and operations of the firm.

THE ACCOUNTABILITY-CAPABILITY GAP

The distance between the responsibility a person holds for AI-influenced oversight, controls and decisions and their actual ability to exercise that responsibility. Accountability without the capability to exercise it is governance theatre. It satisfies the framework, but protects nobody.

The level of AI awareness and capability required will not be the same for everyone. It should be proportionate to the accountability held, the consequence of the decision or control, and the degree of reliance placed on AI. What that means differs materially by role, from Board oversight and challenge to control ownership and the ability to interrogate AI-supported work.

This goes well beyond AI literacy or training completion. The question is what practical capability people need to remain accountable for oversight, controls and decisions increasingly influenced by AI. We return to this in Section 4.

Access and routine use are not the same as capability. Summarising a document or drafting an email demonstrates use; it tells us little about whether someone can use AI effectively in the context of the responsibilities they carry.

For senior leaders, part of that capability is knowing how to use AI to strengthen judgement rather than simply accelerate tasks: interrogating an assumption, testing a decision from a perspective they have not considered, identifying missing evidence or challenging their own reasoning. That requires context as well as access. A general-purpose system knows very little about the person using it, what they are accountable for, or the precedents and experience that shape how they decide.

The uncomfortable possibility is that the people with the most judgement to amplify may be using the smallest proportion of what the technology can do, while the expertise and decision context that would make it useful remain largely in their heads. This is not a question of whether senior people are capable. It is whether organisations have defined and developed the AI capability required for the responsibilities they carry. Boards get briefings, practitioners get tools and training, but accountable leaders need something more.

TWO

It removes some of the occasions on which judgement is formed and tested

As AI takes on more of the work, it removes occasions on which a person would previously have had to form an independent view. Reviewing AI-generated output is a genuine skill, but doing it well depends on knowledge typically built through experience.

Judgement is not built simply by making decisions. It is built by making them, living with them and finding out whether you were right. Form a view, commit to it, see what happens and adjust. Remove those experiences and we risk weakening people's ability to form independent judgement.

The Commission identifies the same challenge. It warns that automation may undermine entry-level training as core processes become black-boxed, and cites declining employment among 22 to 25 year olds in highly AI-exposed occupations while more experienced employment held steady or rose.^{11, 12} The concern is not simply fewer entry-level roles, but the loss of opportunities to learn by doing.

Time to competence means little until firms decide which standard they are measuring: the ability to complete the task with AI, or the ability to form an independent view, challenge the output and remain accountable for the outcome.

A firm can therefore draw down the judgement it already has while creating fewer opportunities to develop the next generation. In the short term, that can look entirely rational: less junior effort, faster output and experienced people reviewing the result. The consequences may only become visible later.

The cost we may be accumulating

We call this capability debt: the gap that can accumulate when AI removes or compresses some of the experiences through which professional judgement develops, without firms deliberately replacing them.

Like financial debt, taking it on may be entirely rational. AI can deliver faster work, lower cost and significant productivity gains today, while the consequences may emerge much later through longer time to competence, weaker independent challenge and fewer opportunities to develop experienced judgement.

The question is whether firms are gaining productivity today at the expense of the capability they will need tomorrow.

THREE

It can erode differentiation

Without organisational context, a general-purpose model draws on patterns learned from a much wider pool of information. If firms use the same handful of models to produce proposals, board papers, credit narratives and analysis, there is a reasonable chance their thinking starts to converge. That does not mean the work gets worse. In many cases, AI may raise the quality floor. The less obvious risk is that it also starts to reduce what makes the work, and the thinking behind it, distinctive.

AI may raise the quality floor while quietly lowering differentiation.

The concern is not just that everyone begins to sound the same. If organisations use the same models with little institutional or individual context, they may increasingly frame problems in similar ways, reach for the same evidence and arrive at similar judgements. Sounding alike is a presentational problem. Thinking alike is a competitive and supervisory one.

For a regulated firm, it also creates an ownership problem. An output that reflects the general case rather than the firm's own methodology, risk appetite and precedent is harder to own, test and evidence. If a suspicious activity report could have been produced by any firm working from the same prompt, it becomes harder to explain why this firm reached that conclusion and who stood behind it. Some institutions are already moving the other way: Morgan Stanley grounds its AI tools in proprietary research rather than leaving them to the general case.¹³

OUR VIEW

General-purpose AI models, given neither firm nor individual context, tend towards more generic outputs. Those outputs may be reasonable in general but much harder to defend in the specific context of a decision. A firm that cannot explain what made an output its own will struggle to show how the person accountable exercised judgement over the decision behind it, precisely what a supervisor, auditor or court may want to understand.

WE OFFER THIS AS A HYPOTHESIS

If organisational distinctiveness is being averaged out, we would expect reduced variance in comparable outputs within firms, greater convergence in problem framing between firms, and a weakening relationship between an output and the context it was meant to reflect. Little of this appears to be measured today. We would welcome the chance to test it with firms, the Commission or academic partners.

FOUR

It widens the gap between people

AI adoption is unlikely to affect everybody equally. Differences in confidence, permission, experience and working norms mean the same technology does not produce the same capability. The Commission finds workforce impacts unlikely to be gender neutral, with women disproportionately in

roles more exposed to earlier task automation.¹⁴ Research using the New York Fed’s Survey of Consumer Expectations found 50 per cent of men had used generative AI against 37 per cent of women, with knowledge differences explaining almost three-quarters of the gap.¹⁵

Gender is the dimension with the clearest evidence, but it is unlikely to be the only dividing line. Neurodiversity, age, disability, role type and seniority are all plausible lines along which adoption may diverge. HM Treasury’s AI Adoption Plan asks that the sector-wide skills plan consider equitable access explicitly.¹⁰

OUR VIEW

A firm reporting 80 per cent active usage has told itself nothing about which 20 per cent are absent, and even less about the 80 per cent, because active can mean a prompt a month or daily use on material work. It conceals who is not building capability, and how uneven that capability already is.

What we learned by doing it ourselves

We have been running an AI transformation programme for around eighteen months, so we have first-hand experience of what happens when a professional services firm sets out to adopt AI seriously.

We do not have all the answers and we are still learning. The Commission has asked the sector what should be done, and evidence from a firm that has experienced some of the challenges first-hand seems more useful than a position.

What we did

From the outset, we made the goal of the programme very clear: trusted, explainable, value-generating AI that enhances, not replaces, the expertise at the heart of our firm. We have repeated that message consistently, and the feedback from our team is part of the reason we keep repeating it.

Prior to the programme our entry point was fairly typical: Copilot, initially for the senior team and then rolled out more widely, followed by ChatGPT, used mainly for sales enablement: research, meeting preparation, proposal drafting and support for other operational processes. Those were the obvious first use cases and they produced obvious benefits, which is why almost every firm starts in the same place.

The use cases expanded faster than we expected. This year we began building structured prompts and reusable skills for client work, which turn our domain expertise and institutional knowledge into an explicit, executable artefact. Guardrails sit on the inputs and outputs, a named person owns the result, and we have tracked adoption since the beginning. In a regulated setting, that engineering, not raw model capability, is what makes AI safe to use on live client work.

On client work, how the skills are built matters as much as what they do.

They are developed and tested with the team, so that people understand how a skill is structured and how it operates. The closest analogy is briefing a colleague: you set out the task, the inputs, steps and the constraints and what a good response looks like.

Regardless of grade, the team owns the output. It is a starting point for analysis, an assessment or a piece of design work, not a finished deliverable, and the person who produced it must review and approve it before it reaches their manager or the engagement lead. That is a deliberate design choice rather than a procedural one. It puts the independent determination, the point of challenge and the named owner into the workflow itself.

What surprised us

Relatively early in the programme we ran our first firm-wide AI adoption survey, structured around a recognised change readiness framework.¹⁶ We report themes here rather than scores, and we would expect several to look different now. Three observations matter for this paper.

Enthusiasm was not the constraint

Appetite to learn was the strongest dimension in the survey by a clear margin, and understanding of why the firm was adopting AI was close behind. Almost everyone believed that building these skills would matter to their career.

The most cited concern was the subject of this paper

Asked about their biggest hesitation, the largest single group named over-reliance and deskilling: loss of critical thinking, analytical skills and human judgement. That theme was cited more often than accuracy, data governance or career security.

Junior colleagues were the most exposed

The most junior grade scored lowest across the survey and was markedly more likely to see AI as a replacement, despite amplification being the programme's message from the start. When AI starts to take on more of the work these are genuine concerns that need to be addressed.

OUR VIEW

We had assumed the constraint would be willingness to adopt, it was not. The team's main hesitation was over-reliance and the loss of judgement, the cost this paper calls capability debt, before we had a name for it

Eighteen months on, the questions the team raises are harder. The most useful challenge was that judgement has not disappeared; it has moved earlier. Designing a structured skill requires critical thinking about which inputs matter, which constraints apply and what a good answer looks like. That embeds professional judgement into the workflow before AI does anything.

What I learned

As Chief Innovation Officer and the person accountable for the programme, these are my own reflections rather than a firm position. In a paper about human accountability, that distinction feels worth making.

People don't move at the same speed. Some of the team went deep very quickly, experimented, innovated and took the initiative to build their own skills and solutions. Others have remained sceptical however much training and reinforcement we have provided. Both groups are useful and I've learned to value both. The early adopters show you what's possible, the sceptics tell you where it breaks and why somebody won't trust it. The mistake is averaging them together and calling the result "adoption".

AI awareness and capability at the top are not optional. I started by getting hands-on, building agents and using AI as a thinking partner with the required context. Not because I had any ambition to become an AI engineer, but because I wanted to understand what these tools could do, where they worked and where they fell over. It has made me far better at judging whether an AI output is genuinely good, because I had been through the process of building, testing and breaking them myself. That doesn't mean every senior leader needs to build agents, but awareness or using AI solely as a productivity tool isn't enough either.

Involving the team in building the skills has been the best training we have found. It forces experienced people to articulate how they approach a task, often for the first time. It teaches the team what the skill can and cannot do, creates ownership because the output remains theirs to stand behind, and spreads capability rather than concentrating it in the handful of people who naturally went deepest first. That has probably changed my view of AI adoption more than anything else we have done.

Measuring capability uplift is challenging. Licences issued, use cases live and skills deployed tell you very little. What I want to know is whether people are becoming more capable, more confident and better at challenging what comes back. We have limited ways of measuring that systematically and are actively working on it. It is also the argument of this paper, and it would be disingenuous to make that argument without admitting we've not fully solved it ourselves.

What choosing Human-Led AI requires

Four actions follow, addressing the awareness and capability firms need, how to measure and enable it as AI does more of the work.

1 Equip the people who carry accountability

DEFINE ACCOUNTABLE AI CAPABILITY BY ROLE AND TEST WHETHER IT EXISTS

The AI awareness and capability required should be proportionate to the accountability held, the consequence of the decision or control, and the degree of reliance on AI. This does not require universal technical expertise. A Board member does not need to interrogate individual AI outputs, but they do need to understand where AI materially influences the business, challenge what the firm relies on it to do and judge whether governance and assurance remain sufficient.

ROLE	WHAT GOOD LOOKS LIKE	PRACTICAL AI CAPABILITY
Board	Can challenge where AI is materially used, what the firm relies on it to do, what could go wrong, how management knows it is working and whether governance and assurance are sufficient.	Has sufficient practical experience of AI to understand its capabilities and limitations. Can use it to interrogate assumptions, explore alternative perspectives and challenge evidence, while understanding how context shapes the quality of its output.
Executive / accountable senior leader	Can make informed decisions about deployment, risk and intervention, and recognise when adoption is moving faster than governance, controls or capability.	Uses AI as a thinking and challenge partner: testing assumptions, exploring alternatives and identifying missing evidence, while providing the organisational and decision context needed to make outputs relevant.
Control owner	Understands how AI changes the design and operation of the control, where human judgement remains critical, what failure looks like and when intervention is required.	Can contextualise AI to the firm's methodology, risk appetite and control environment, interrogate outputs and determine where AI can and cannot be relied upon.
Manager / reviewer	Can interrogate AI-supported work, identify inappropriate reliance, challenge the evidence and determine when additional human review is required.	Can use AI to test reasoning and identify gaps while independently assessing quality; understands how prompts, context and source material influence outputs.
Practitioner	Can use AI appropriately, understand its boundaries, interrogate outputs, recognise plausible errors and challenge, override or escalate when required.	Can give AI sufficient task and domain context, interrogate rather than simply accept its output, refine its reasoning and apply independent professional judgement to the final result.

WHAT THIS MEANS IN PRACTICE

- Treat AI capability as role-specific: Board, accountable executive, control owner, manager and practitioner need different standards.
- Test whether people can challenge AI outputs in their own domain, rather than inferring capability from training completion.
- Start with the roles carrying the most accountability and the most reliance on AI.

2 Change what you measure

EVIDENCE CAPABILITY, QUALITY AND DISTRIBUTION

Stop treating licence deployment as evidence of adoption or capability. Firms need to look beyond access and activity: can people evaluate AI-supported work in their own domain and explain why they accepted, challenged or rejected it; is the work improving in quality as well as efficiency; and who is not adopting, or is using AI only superficially, and why?

MEASURE NOW	EXPERIMENT NEXT	STANDARDISE LATER
Most firms already hold richer adoption data than they report. Break it down by role, seniority, team, frequency and depth of use, and the diversity dimensions the firm already tracks.	Develop and test measures of quality, challenge, time to competence, the ability to spot plausible AI errors, control effectiveness and the capability of people carrying accountability. None of them is settled.	Only once measures are proven feasible inside firms should the sector try to make them comparable between firms. Starting with comparability is how measurement programmes stall.

WHAT THIS MEANS IN PRACTICE

- Look beyond licences and headline usage to understand the depth and distribution of adoption.
- Track quality, rework and review findings alongside efficiency and productivity.
- Experiment with measures of capability and time to competence rather than claiming they are already solved.

3 Give AI the context it needs

MAKE INSTITUTIONAL KNOWLEDGE USABLE AT THE POINT OF WORK

Put the firm's methodology, standards, precedents and institutional knowledge into the systems people actually use, rather than leaving them in documents beside the work. Treat that context as infrastructure, jointly owned by those responsible for the knowledge and the technology, and test whether AI-supported outputs reflect the firm's methodology and risk appetite rather than the general case. Without it, firms risk using the same technology to arrive at increasingly similar answers.

4

Redesign how judgement is formed

REBUILD THE LEARNING LOOP AS AI CHANGES THE WORK

Do not preserve the old apprenticeship for its own sake. Identify what junior work actually taught people, separate that learning from the low-value effort around it, and deliberately recreate the experiences that developed judgement. Where judgement matters, people should still have opportunities to form a view, test it against AI and find out whether they were right.

Keeping a human in the loop is not enough. A click-to-approve workflow and an expert independently challenging an AI-supported conclusion are not equivalent controls. For material decisions, firms should be clear about where human judgement matters, what challenge is expected and what evidence should be retained.

Assurance should examine the human, the AI and the process together, not simply whether a human was present.

WHAT THIS MEANS IN PRACTICE

- Ask teams to form a first view before using AI on selected tasks.
- For material decisions, define the human challenge expected and retain evidence of how the final decision was reached.

OUR VIEW

Workforce capability and AI assurance are often treated as separate subjects owned by separate functions, they are parts of the same control system. A capable workforce is a control; an incapable one is a weakness that model governance alone cannot address.

Individual firms can act now, but some of these questions will ultimately require a collective response. Common expectations for AI capability, how professional judgement is developed as work changes, and how human and AI contribution should be assured will need to evolve across firms, regulators, professional bodies and government. The priority now should be to test what works in practice rather than standardise too early.

The question worth leaving with

It would be easy to read this paper as an argument for caution, but it is not. AI will do more of the work, and firms that resist that shift will lose to those that embrace it; the challenge is to make sure that, as the technology becomes more capable, the human system around it is designed just as deliberately.

The aim is not to preserve human effort. It is to amplify and increase human capability.

Done well, AI could help people reach independent judgement faster than previous generations, not by preserving the old routes through which experience was built, but by deliberately designing better ones around the work people will actually do in an AI-enabled organisation.

Most firms can already tell you how many of their people have access to AI, how often they use it and, increasingly, what time or cost it saves. As agentic AI moves from helping people perform work to performing more of the work itself, these questions will become more important still, and that is a subject for a future paper.

That leaves the more immediate question:

Can we evidence that our people are becoming more capable as AI becomes more capable?

Few firms can answer that confidently today, which is understandable given how quickly the technology and the work around it are changing. But as AI takes on more of the work, firms will need to become much more deliberate about the awareness and capability they expect people to develop, how that differs according to role and accountability, and how they know it is actually being built.

About this paper

Plenitude is a Supporter of the Financial Services Skills Compact. This paper builds on the Financial Services Skills Commission's research for HM Treasury and is offered to the wider discussion about what AI adoption requires of the workforce in regulated financial services.

AUTHOR

Alan Paterson

Founder and Chief Innovation Officer
Plenitude Consulting

WITH CONTRIBUTIONS FROM

Louise Brett

Chief Executive Officer
Plenitude Consulting

Mark Humphries

Chair
Plenitude Consulting

Sources

1. Financial Services Skills Commission, *A Workforce Transformed: Technology, skills and the future of work in financial services*, May 2026, key findings.
financialservicesskills.org/wp-content/uploads/2026/05/AI-disruptive-technology-report-workforce-transformed.pdf
2. FSSC, *A Workforce Transformed*, May 2026, section 5, "Key challenges".
3. FSSC, *A Workforce Transformed*, May 2026, key finding 9, on the number of highly skilled people the sector must recruit and train by 2035 to replace those leaving or retiring.
4. FSSC, *A Workforce Transformed*, May 2026, on the proportion of tasks in most UK financial services roles expected to change through automation.
5. Bank of England and Financial Conduct Authority, *Artificial intelligence in UK financial services: 2024 survey*, November 2024, including the proportion of firms reporting complete understanding of the AI they use.
bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024
6. FSSC, *A Workforce Transformed*, May 2026, section 1, on adoption, licensing and benefits realisation.
7. Financial Services Skills Compact, 2026: four employer commitments, including annual public reporting of progress. Separate from, though commissioned alongside, the HM Treasury research programme. financialservicesskills.org
8. FSSC, *A Workforce Transformed*, May 2026, introduction, on the scope of the HM Treasury research programme. Phase One assessed how technology is changing skills and the workforce; the remaining phases report in January or February 2027.
9. Financial Conduct Authority, *The Mills Review: AI and the future of retail financial services to 2030*, 6 July 2026, on the autonomy spectrum and accountability.
fca.org.uk/publications
10. HM Treasury, *Financial Services AI Adoption Plan*, 14 July 2026, recommendations on AI skills, including a sector-wide plan covering boards, specialists and frontline teams.
gov.uk/government/publications
11. FSSC, *A Workforce Transformed*, May 2026, section 4, on entry-level roles and institutional knowledge.
12. Brynjolfsson, E., Chandar, B. and Chen, R., *Canaries in the Coal Mine? Six Facts about Recent Employment Effects of Artificial Intelligence*, Stanford Digital Economy Lab, August 2025, finding a 13 to 16 per cent relative decline among 22 to 25 year olds in highly AI-exposed occupations. Cited in FSSC, section 4. digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence
13. Morgan Stanley, *AI @ Morgan Stanley Assistant and Debrief*, built with OpenAI to ground AI tools in the firm's proprietary research library. morganstanley.com/what-we-do/technology
14. FSSC, *A Workforce Transformed*, May 2026, on the distribution of workforce impacts.
15. Aldasoro, I., Armantier, O., Doerr, S., Gambacorta, L. and Oliviero, T., *The Gen AI Gender Gap*, BIS Working Papers No. 1197, July 2024. bis.org/publ/work1197.htm
16. Plenitude, firm-wide AI adoption survey, 2026. Themes reported, not underlying scores.

Plenitude

Built to combat the threat, Plenitude is a specialist economic crime firm driven by experience and accelerated by AI.

We work across every stage of the customer lifecycle: advising, transforming, assuring and empowering clients to combat economic crime, meet regulatory obligations and reduce the harm it does to society. Appointed to the FCA Skilled Persons panel for Financial Crime, Market Abuse and Conduct, we help clients detect, prevent and remediate these threats with confidence.

We have a 14-year delivery track record, supporting some of the most challenging programmes in the industry across global banks, insurers and asset managers.

THE AI LAYER

We engineer trust from the outset, not as a by-product. Assurance holds the standard, the delivery model runs on every engagement, and our RegTech and AI solutions drive more effective risk management.

01

AI Assurance Framework

Adapted from the NIST AI Risk Management Framework, operationalised through our Risk Taxonomy and G3MA Assurance Framework (Govern, Map, Measure, Manage, Assure).

02

AI-enabled delivery model

On every engagement. Engineered, not improvised: persistent context grounded in Plenitude methodology, engineered guardrails, and a named person accountable for every output. Days compressed to hours.

03

RegTech & AI solutions

An 8-year track record building and deploying RegTech products and AI solutions used by some of the world's largest financial institutions: Supernova, RegSight and Compass.



plenitudeconsulting.com · LinkedIn: Plenitude Consulting

ai@plenitudeconsulting.com