When the tech works but the organisation doesn't

What FS&I leaders are saying about AI, two years in.

Drawn from a Chatham House discussion hosted by SCC with Dr Leda Glyptis, London, October 2026. For CIOs, CTOs, CISOs and heads of data and AI in financial services and insurance.



Two years in, and not much to show for it.

Financial services and insurance has been doing AI for two years. The tools work and people enjoy using them; many leaders say their teams have had more fun at work in that time than they can remember. And yet the bills are high, the operational gains are hard to find and the risk picture has got worse. The press releases say everyone is ahead. In private, the picture is different.

In October 2026 we invited a group of financial services and insurance leaders to spend an afternoon discussing the state of AI in their industry, together with Dr Leda Glyptis, who opened the afternoon and sat in on the discussion that followed. Leda is Visiting Professor of Practice at Loughborough University London, a non-executive director, adviser and author, and has spent 25 years in and around financial services technology, including inside the banks now trying to make AI pay off.

Source: The AI Process Gap, September 2026

72%

of AI initiatives fail or stall

on process-related challenges, at an average cost of $1.55m per organisation.



A tool, not a transformational moment.

Leda's view

Her view is that the sector approached the capability as a tool rather than as a transformational moment, then compounded the miss twice over. It treated the tool as a silver bullet that would single-handedly fix everything. And it assumed that the moment it chose to engage, the technology was fully formed and complete, when it was and remains a moving target. It spent accordingly.

The money went in and the operational value has not come out: firms either cannot quantify the gain or report single-digit uplift, and a September 2026 study she cited found 72% of AI initiatives fail or stall on process-related challenges, at an average cost of $1.55 million per organisation (The AI Process Gap, 2026). The AI did what it was asked. What it was fed was too inconsistent for it to do anything useful. Meanwhile token fraud, self-repairing attacks and rogue agents have arrived and regulation has not. Her summary: expensive, increasingly sinister and not much to show for it, while employees want more.

> "Calm, but not caution. Do not slow down. Do not back down. Make the decisions deliberately."
>
> **Dr Leda Glyptis,** closing her keynote



Dr Leda Glyptis

Visiting Professor of Practice,
Loughborough University London
Non-executive director, adviser and author

Six questions Leda put to the room.

01

What does governance look like in a world of plenty?

Governance and prioritisation have always been linked by scarcity: this many engineers, this many hours, so this gets done and that does not. When anything can be built, scarcity no longer sets focus. Something else has to. Does governance sit with who has the tools, or with what they are allowed to do? Does the risk register need rewriting?

02

What is the operating model for an AI-enabled business?

If the technology is transforming the market and the workforce, which parts of the current operating model need to change, in what order, and what breaks first? Leda has never seen a properly funded target operating model in her career. The technology gets funded. The model it lands in does not.

03

Can we carry the debris with us?

Every organisation holds silos nobody has broken down, data nobody has harmonised, policies that are no longer fit, workarounds accepted as risk and customer journeys designed around what the systems could do rather than what the business was trying to achieve. Nobody wants to fix it. Can the sector afford not to?

04

What is a valuable skill set now?

The newly appointed chief AI officers she has met are long-serving executives with a new title. They know the organisation. They are not AI specialists. A good engineer today has a different profile from a good engineer five years ago, and the apprenticeship model used to train them has not caught up.

05

What is procurement's job?

Procurement is the most neglected department in the building and is about to become one of the most important. How a contract with a model vendor was negotiated is a measure of the third-party risk the firm now carries. Roll-back, portability and concentration risk belong in that contract. Procurement was not in the room for this conversation, and in Leda's experience it rarely is.

06

Is cyber risk still a bookend?

Security used to be assessed at the start of a programme and tested when circumstances changed. Now safety, resilience and the speed of change are one conversation. Automated attacks can only be detected by automation, so risk and compliance teams are asking for AI tools to defend against AI. Procurement and cyber need to work as a single function, possibly as a single role.

Govern with frameworks, not rules.

The technology changes faster than any rule can be written. By the time a control is signed off, a vendor has shipped something that makes it irrelevant. Governance has to be framework-based: principles people are held to, with a named business outcome behind each control. Few firms can state that outcome today, and the way AI has been rolled out so far let corners be cut and exposed existing governance as inadequate.

In practice, governance is local. It follows business unit structures, so each unit has its own tools and its own rules, and nobody has settled how much standardisation is right or who owns what. Where it is working, it is governance by resource: token controls at a high level, personal scorecards for code quality and usage, and the same pre-release gates as before AI. Budgets get set by letting a small group use the tools, setting soft targets from what you see, then moving to hard limits. There is no way to set the target up front.

Two risks sit underneath. Token cost and compute capacity will hit a ceiling, and nobody can see when, or how to avoid becoming hostage to fortune. And a regulator asked why it did not intervene in a recent cloud outage may act on AI before firms are ready, possibly with a senior managers regime that makes a named individual accountable for explaining the logic of an automated process. Nobody in the room liked that idea. Nobody ruled it out.



Nothing is new, except the speed.

The strongest sentiment in the room was against treating AI as an exception. People used to pull data into spreadsheets they should not have. Now they write code. The legislation protecting customers and data, the obligations on the firm and the principles are the same as for any cloud service. What has changed is how fast the technology arrived and how fast things unravel when they go wrong.

The shadow AI is in the platforms you already pay for

Speed is also why shadow AI is a bigger problem than firms think. Few can say for certain what AI is running inside them, and the gap is widest in the platforms they already pay for. A vendor adds an AI feature to a long-standing tool, switches it on for one team to see if it works, and tells nobody in procurement, security or technology, because none of them manage it. The fix is technical controls that can see what is running, with cyber posture driving procurement decisions, rather than more people checking. Annual click-through policy training will not close the gap, and nobody announces a new policy in a way people remember.

The skills are already eroding.

Coding and debugging skills are eroding faster than anyone expected. One example cut through: engineers telling their manager they could not work because they had run out of tokens. They could code before tokens existed. Roles are blurring at the same time. Business analysts, whose job was the problem statement, now build things themselves, and developers build well but struggle with the problem statement.

That matters because humans still own assurance. The firms with this under control test heavily, then have code reviewed by a less capable model than the one that wrote it, using several reviewer personas, with failures sent to a person. Use the same model to write and review and errors go undetected and unravel faster. The skill that is eroding is the one needed to check the thing that is eroding it.

The direction of travel is engineers moving into analysis and product work, T-shaped specialists with real depth and a premium on collaboration. That flows into hiring, incentives and retention, and few firms have adjusted any of them. No skill set becomes redundant, though at least one firm has let a department go to test that view.

> “The skill that is eroding is the one needed to check the thing that is eroding it.”
>
> **Dr Leda Glyptis**

Stop being vertical.

The business is siloed, processes and data are all over the place, and the operating model is not ready for AI. Centralising everything and starting again runs into HR, legality and reality. The useful move is to spend less effort documenting existing processes and more on outputs: what the business is trying to achieve for customers, within what guardrails and where it needs people. People stay in the model for judgment and emotional intelligence, less for transactional work.

The AI transformations that are working sit in fluid, purpose-led environments where multidisciplinary teams already exist. Where the structure is vertical, the technology lands on it and slides off.

Data is the same problem. Different departments will always hold different data; what breaks things is inconsistency inside each set, the same field entered and labelled differently by different people. Fix that at source and the pieces assemble. The everyday version is the salesperson who leaves the CRM wrong because his manager asks for the same numbers in a spreadsheet, then a slide deck, then on a call. Change how data is reported upwards and the behaviour follows.

> “Where the structure is vertical, the technology lands on it and slides off.”
>
> **Dr Leda Glyptis**



We are running out of technology deliverables.

Financial services has spent 20 years delivering technology programmes, and delivering them well. Each one touched policies, governance, ways of working and other departments, and in nine cases out of ten the budget did not stretch to the work around the edges. So the programme logged a risk acceptance on the compliance process, agreed a workaround where it handed over to another department and recorded an exception on procurement. After two decades, that layer of accepted exceptions is dense.

The big transformations are now running out of technology deliverables. What remains is the connective tissue: deep dependencies, almost no documentation and very little appetite, because a chief executive who takes it on is staking their reputation on cleaning up a mess. Some of it may not need doing at all.

Teams are still producing documentation for regulatory requirements that lapsed years ago because nobody remembered to stop them. Whole chunks of activity can be bypassed rather than modernised, but you need to do the difficult work to find out which.

The same logic applies to what firms ask of AI. The request is usually to automate the process as it is. ‘Do not modernise’ the process. Imagine it does not exist and start again, or the firm ends up using AI to rebuild a spreadsheet one person built 20 years ago, then using AI again to get the data back out. Every firm has that person: someone who built the year-end process or the exotic-currency holiday macro, never documented it and has since left. Point an agent at the raw data and it has none of that context. The connective tissue runs between people as much as systems, and it leaves when they do.

The question that split the room:

are large firms doing real transformation, or bolting AI onto what exists and calling it transformation?



Why it stays undone.

The work is known. The reasons it does not get done are human.

The first is that it is treated as a technical problem when it is social and political. AI makes data transparent up and down the organisation and joins processes across functional boundaries, which cuts directly across embedded roles, power structures and politics. Middle managers and specialists spend much of their time arbitrating between the coalface and the strategy, none of it written down. A job description is mostly papering over cracks, and the people who understand how things really work are being asked to make themselves redundant. They know it.

The second is the position of the chief executive. Data is only accurate now. Asking for £150 million to rebuild the data architecture so the AI programme can work competes with every other claim on the budget, and with the fear of being pilloried for the choice a year later. Ask chief executives what they focus on and the answer is making money and staying in a job. Not deciding serves both for a while. Everything discussed in the room plays out over the next couple of years and will change how firms make money, so the trade-offs are coming regardless.

Culture and organisational legacy take longer to shift than any technical deployment. That is the argument for starting now.

> “Culture and organisational legacy take longer to shift than any technical deployment. That is the argument for starting now.”

Dr Leda Glyptis



Six questions to take to your own leadership team.

Every conversation in the room came back to the same pattern: the technology was ready, the organisation around it was not. These are the questions that exposed the gap. Ask them of your own executive team and see how many get a confident answer.

01

Governance: what is our governance for?

Does it exist to control who has the tools, or what they can do with them? Can anyone name the business outcome each rule protects? And when engineering time is no longer the constraint on what gets built, what is? If prioritisation still runs on scarcity that no longer exists, governance is managing the wrong thing.

02

Operating model: have we funded the change, or just the technology?

Which parts of the current operating model stop working in an AI-enabled business, in what order, and who owns redesigning them? Few firms could point to a funded plan for the operating model itself. Many could point to a funded plan for the tools.

03

The debris: what have we been carrying for twenty years?

Every programme leaves behind risk acceptances, workarounds and exceptions logged because the project did not touch them. Some teams are still producing paperwork for regulations that lapsed years ago. Do we know where ours is, and which of it we can bypass rather than modernise?

04

Skills: do we know what good looks like now?

A good engineer today is not the engineer of five years ago. Are we still training and hiring for the old profile? What skills are eroding because people no longer use them, and which do we need to protect? And does the person with “AI” in their title know the answer, or just know the organisation?

05

Procurement: could we list every AI capability running inside a platform we already pay for?

Vendors switch features on and enable them for a single team. If nobody tells procurement, security or technology, the firm has shadow AI it did not choose. What do our contracts say about roll-back, portability and what the vendor can do with our data? Who in procurement has read them with that in mind?

06

Cyber: are we quick enough?

The obligations have not changed. The speed at which things go wrong has. Can our security team detect an automated attack at machine speed, or only after the fact? And do cyber and procurement work as one conversation, or two departments who meet at contract renewal?

Next steps

Answer the six questions honestly

Book a 60-minute conversation with SCC.

We will work through the six questions with you: which ones your firm can answer today, which ones go quiet and what it would take to answer them honestly. You leave with a clear view of where you are and a first piece of work to start on.

Book Now

Speak to your SCC account manager or scan to book