



Paper 1

AI Through the Eyes of 40 Delivery Professionals

What's really holding adoption back in insurance



The Reality of AI in Insurance

Most organisations are further behind than they will admit – and further ahead than they give themselves credit for.

That tension came through clearly when we brought over 40 of our consultants together for the first of our AI-focused Communities of Practice sessions.

The question we put to the room was straightforward: how is AI actually changing the way you deliver?

What we got back was not a room full of AI converts talking about transformation. It was something considerably more useful: an honest account of where things actually stand, from Business Analysts to Delivery Leads working across the insurance market every day.

This paper is our attempt to share that picture. It is not Stellarmann's view. It is what 40 professionals said when they got in a room together.



“

Stellarmann has built a powerful network of IT and change experts across insurance and financial services. Our Communities of Practice invest in that expertise – creating a space to share hard-won insights from the front line, challenge what works, and yes, unwind after major projects. The goal is simple: surface what actually delivers, and take it straight back to our clients. In a market moving this fast, honest, peer-to-peer insight beats any vendor pitch.

Matt Coggan, Stellarmann



The Experiment is Over. Now Comes the Hard Part

Something shifted in 2025 and into 2026 – and if you are working in insurance, you will have felt it.

For the two years before that, AI adoption was largely characterised by experimentation. Innovation labs explored use cases. Technology teams ran proofs of concept. There was considerable investment in demonstrating what AI could do. As Scott Davies, one of our practice leads who works closely with insurers and financial services clients, observed:



Too often, these failed to become adopted and operational – because they were more technology-centric, assessing technical feasibility in isolation from the business.

Scott Davies, Stellarmann



That period of experimentation is drawing to a close.

AI is now built into the core platforms insurers rely on – policy administration, claims, CRM, and ERP. It's no longer a side initiative; it's woven directly into the systems underwriters, claims handlers, and risk managers use every day.

That changes the nature of the challenge considerably.



The question is no longer "what can AI do?" It's become "How do we govern, scale and extract genuine business value from what is already here?"

The organisations answering that question well are those that shifted early from a technology-centric mindset to a decision-centric one. They identified the decisions that matter most – high-value underwriting calls, complex claims assessments, risk pricing – and asked how AI could make those decisions faster, more consistent and more defensible. Everything else followed from that clarity of purpose.

That shift in thinking also demands a shift in how programmes are delivered.

Many of the most successful organisations have moved from centralised technology projects to a product model – small, business-owned squads focused on specific decisions or value streams, with AI engineers embedded alongside underwriters, claims handlers, and risk managers. They're not building AI for the business; they're building it with it. The result is real adoption, because the people who use it help shape it from day one.



5 Important Conversations the Industry is Avoiding

1

Data Foundations Aren't as Solid as Leaders Think

AI in insurance lives or dies on data quality. Underwriting decisions, claims assessments, risk pricing – all depend on data that is accurate, consistent, complete and accessible. Most insurers have meaningful gaps in at least one of those areas, and some have gaps in all of them.

The focus across AI initiatives in our client base through 2025 has been consistent: better data in, better decisions out.

That means building structured, governed data environments – what some clients refer to as a Data Universe – through which internal and external data flows in a way that AI can actually use. It means understanding data lineage, establishing clear ownership and mapping precisely how information moves between policy, claims and finance systems.

Without that foundation, AI produces outputs that are inconsistent, difficult to explain and impossible to scale responsibly. In a regulated industry where every significant decision requires justification, that is not a sustainable position.



2

Governance Arrives Too Late to Protect Decision-Making

The technology is outpacing the governance structures designed to manage it – and in insurance, where regulatory expectations are high and the consequences of poor decisions are material, that represents a genuine and growing risk.

A significant number of organisations are still treating governance as something to establish once an AI model is live.

The firms making the most consistent progress are taking the opposite approach: building the necessary guardrails – around data use, model explainability, audit trails and human oversight – before deployment, not as an afterthought to it.

In the Lloyd's market and among London Market carriers, where underwriting authorities and delegated authority arrangements create specific and well-defined governance obligations, this sequencing matters considerably. An AI model that cannot account for how it reached a particular decision is not only a technical limitation. In many cases, it is a regulatory exposure.



3

Outcome Assurance is the Weakest Link

Ask most AI programme leads whether their model is producing the right outputs and you will receive a confident answer. Ask them how they know – what the assurance mechanism is, how they would identify and respond to model drift, how they would account for a specific decision if asked by a regulator or a claimant – and the picture becomes considerably more complex.

This is the most significant barrier to AI being deployed in genuinely high-value decisioning across insurance.

Until organisations can answer those questions with confidence, AI tends to remain confined to lower-risk activities: document summarisation, data extraction, internal search tools. Valuable in their own right, but far short of the potential.

The organisations that are unlocking real value – using AI to inform underwriting decisions, support claims triage, improve exposure management – are those that have invested in robust outcome assurance frameworks. They measure model performance on an ongoing basis, build explainability in from the design stage and maintain clear escalation pathways when AI output is questioned.



4

Most AI Programmes Fail on the Journey from Pilot to Scale

A successful proof of concept and a scalable operational product are two fundamentally different things – and the distance between them is frequently underestimated.

Business ownership is often unclear. Edge cases have not been stress-tested. And when the time comes to transition from pilot to product, the gap between what was built and what the business actually needs is often greater than anticipated.

That gap is compounded by the state of the tooling itself. While AI capability is advancing rapidly, many of the enterprise platforms teams rely on day-to-day have been slower to evolve.

Keren Klein, an AI Transformation Coach specialising in embedding AI agents into enterprise workflows reflected:



Existing tools have been slow to catch up with the pace of AI development. When AI features arrive, they're often restrictive, which forces teams to work around limitations rather than fully embedding AI into how they operate.



Even where progress is being made, adoption remains uneven. Tools such as Miro and LeanIX have introduced AI capabilities, but often later than expected and with constraints that limit integration into real workflows. The result is a gap between what AI can do and what can be operationalised at scale.



5

AI Adoption Stalls Where Trust Is Assumed, Not Designed

Much of the AI discussion in insurance focuses on capability, governance and outcomes. Far less attention is paid to a quieter, but equally material factor: whether people actually trust AI enough to use it.

Across the room, consultants described a familiar pattern. AI tools were available and the value was understood, yet adoption remained hesitant. As [Philippa Rocks](#), Change Manager within a major transformation programme, puts it:



Adoption doesn't fail because the AI isn't capable. It fails because people don't trust it yet. There's an assumption that if the business is engaged, adoption just follows - but in reality, trust has to be earned.

That experience mirrors wider public sentiment. Research by the [Tony Blair Institute for Global Change](#), conducted with Ipsos, shows that UK adults are significantly more likely to view AI as a risk to the economy than an opportunity, with lack of trust cited as one of the most frequently identified barriers to adoption.

Trust comes from evidence, not advocacy. Without visibility into how decisions are made, adoption doesn't follow.

Keren Klein also emphasised the importance of trust.



Even when teams recognise the value AI adds, adoption is slow. People want to understand how the AI is thinking before they're comfortable relying on it, especially when the output affects their work.



Behavioural factors amplify this challenge. Age and career stage were repeatedly cited as influencing openness to AI, with younger colleagues generally more willing to experiment and more experienced staff often cautious – driven less by capability gaps and more by concerns around accountability, redundancy or disruption to proven ways of working.

These dynamics are compounded by tooling limitations. While AI capability is advancing rapidly, many enterprise platforms have been slow to adapt, offering constrained integrations that make meaningful adoption harder than it needs to be.

Taken together, these insights point to a consistent conclusion: AI adoption is not simply a technology or delivery challenge. It is a trust challenge. Organisations that fail to address the emotional and behavioural response to AI – fear, scepticism and perceived loss of control – will continue to see adoption plateau, even where underlying capability is strong.



What Successful Delivery Looks Like on the Ground

The organisations making consistent, sustainable progress with AI in insurance are not necessarily those with the most advanced technology or the largest budgets.

What differentiates them is not capability, but how they deliver and how they own the outcome. Across our client work, a clear pattern has emerged...

Successful organisations have moved decisively away from traditional, centralised project structures.

Rather than treating AI as a discrete technology programme owned by IT, they operate through multiple business-owned delivery squads aligned to core functions such as underwriting, claims and pricing. Each squad is accountable for a specific decision set or value stream and is measured on outcomes rather than activity.

AI engineers work directly alongside underwriters, claims handlers, actuaries and product owners. The focus is not on building models in isolation and handing them over, but on solving live business problems together. Where that proximity exists, adoption follows because the people expected to use the capability have helped shape it from the outset.



The winners are those who ditch traditional project roles. By moving to product-focused delivery - where AI engineers and business leaders co-own outcomes - teams can test, learn, and iterate quickly, driving continuous customer impact.

Scott Davies, Stellarmann

These organisations also operate with a fundamentally different delivery rhythm. They are not working towards a single go-live date. Instead, they release incrementally, measure impact continuously and adjust based on real-world outcomes.

Success is defined not by whether the AI model functions technically, but by whether it improves decision quality, consistency and confidence over time.



Questions Leaders Should be Asking Now

For those of you leading or sponsoring AI programmes, the following questions are worth reflecting on honestly:

Data

Do you have a clear, shared understanding of your data – where it lives, how it moves between systems, and who truly owns its quality and use?

Governance

The governance and regulatory expectations built into your AI initiatives from the outset, or being addressed only once models are live?

Outcome

How do you know your AI is producing the right outcomes over time – and could you evidence, explain and defend a specific AI-assisted decision if challenged?

Scale

Are your AI initiatives explicitly designed for operational adoption and long-term ownership, or primarily to demonstrate feasibility and win approval?

Trust

Do the people expected to use AI genuinely trust it enough to rely on it for decisions they remain accountable for – and have you designed explicitly for that trust to develop?

There are no right or wrong answers.

But the organisations prepared to confront these questions directly – across technology, delivery and behaviour are consistently the ones making meaningful headway.





From the Room to the Real World

Everything in this paper came from practitioners working inside live AI programmes in insurance.

It reflects where things actually stand – where progress is being made, where it is stalling, and where delivery reality is challenging the narrative.

The natural next step is to move from insight to application.

Paper 2 takes the themes from this paper and shows how they play out in practice – through real case studies drawn from live insurance environments.

This is not theory or best practice. It is what happened when organisations moved beyond proof of concept and attempted to make AI work in the real world.

[Access Paper 2: Real AI Programmes in Insurance](#)





Continue the Conversation

If any of the themes in this paper resonate - whether that's scaling beyond pilot, building trust in AI-supported decisions, or embedding AI into delivery - we would welcome a conversation.

Joe is leading AI discussions at Stellarmann and works closely with insurers on these challenges. He'd be happy to share his perspective or discuss how they're playing out in your environment.



Joe Pryke

Client Principal

Email: joe.pryke@stellarmann.com

Join our Communities of Practice

This paper is one output from our Communities of Practice - a network of delivery professionals sharing insight from live programmes.

If you're a Stellarmann contractor or delivery specialist and would like to get involved, we'd love to hear from you.

Get in touch with Katie Thomas: katie.thomas@stellarmann.com