Artificial Intelligence

How the insurance industry and actuarial modellers are embracing AI

June 2026

In collaboration with

InsuranceERM

FIS

Introduction

AI has been billed as the next great technological revolution and in the realm of insurance it promises to deliver efficiency gains, improved products and customer service.

But as with any novel development, there is a period of sometimes painful evolution to go through first. We have certainly seen that with AI. Early iterations of large language models such as ChatGPT became notorious for giving dangerous advice and having hallucinations. These errors were noticed, the engineers went back to fix it, and we tried again – and it was better.

In a competitive and highly regulated industry such as insurance, this has created tensions. Those first and second attempts at employing AI might never pass muster with firms and their supervisors. But the rapid pace of AI development means that insurers, if they are to remain competitive, need to keep experimenting and finding potential applications, while designing guardrails to ensure reputations are maintained.

In this special report, we examine what some of the leading insurers are saying about their use of AI (page 2) and what impact it is having on their strategy. We also uncover some case studies showing what they are actually doing with AI (page 6).

Successful use of AI involves careful management of the implementation process (page 11) and managing the risks around AI (page 17).

In actuarial modelling, machine learning has been used for many years and now actuaries are looking to take a step forward with AI, both as a modelling technique and to help with the processes that surround actuarial modelling. We investigate how AI is creeping into actuarial processes (page 13) and how firms can embrace AI without sacrificing governance and resilience (page 10). The complexity of AI models means it's sometimes difficult to explain how they work, so we also examine the approaches actuaries are using to gain confidence in AI models (page 15).

Christopher Cundy
Editor
InsuranceERM

Contents

What the world's leading insurers are saying about AI

Insurers are moving AI from pilot projects to core processes in their business, while trying to maintain awareness of the risks associated with the technology. Martin Assmann reports

By late 2022, it looked like the tech industry was out of ideas. The iPhone launch was over a decade ago, and the internet was into its middle age and truly not exciting anymore.

The time of genuinely life-changing technological breakthroughs had seemingly gone by, replaced by a sludge of incremental changes.

Then OpenAI released ChatGPT 3.5 and the rest, as they say, is history.

Now, three and a half years later, the large language models underpinning chatbots and AI agents have changed the world. The nascent technology has gone from interesting to indispensable for companies wanting to compete – even in the traditionally slow-moving insurance sector.

Insurers are now under pressure to take AI seriously, find practical uses for it and prove the technology can deliver more than just polished demos and buzzwords.

But how far has the transformation actually gone and where is it heading? *InsuranceERM* spoke to senior executives at some of the world's largest re/insurers to assess how they view AI, where it is already being used and whether the promised benefits are showing up in practice.

Not an experiment anymore

At the top of the market, the language on AI has changed over the last 12 months.

Last year, much of the industry spoke about pilots and understanding the technology. Now, insurance executives like Matthieu Caillat, Axa's chief technology and AI officer, tell *InsuranceERM* the technology has made its way to "the core of what we want to achieve".

However, he is also clear about the challenges: "We are going to have to question ourselves about what we use it for, what the methodology is and how we onboard our people".

Chubb chairman and chief executive Evan Greenberg similarly writes in his 2025 shareholder letter that the group is "embracing the power of technology, AI, data and business process change, transforming ourselves so we thrive for years to come".

Ermir Qeli, head of data science and AI, group data services at Swiss Re, tells *InsuranceERM* the reinsurer's thinking has shifted from discrete use cases to process redesign.

> "We are going to have to question ourselves about what we use it for, what the methodology is and how we onboard our people"
>
> **Matthieu Caillat, Axa**

Qeli says he "would describe what we did over the last 12 months as a reimagination of the processes end to end, thinking through what the core role of AI is and where we can transform the way we do our work."

José Luis Bernal, transformation director at Spanish insurer Mapfre offers a more cautious assessment of the industry's progress.

"[I] would describe what [Swiss Re] did over the last 12 months as a reimagination of the processes end to end"

Ermir Qeli, Swiss Re

Speaking to *InsuranceERM*, he rejects the idea that insurers are already close to becoming AI-native businesses.

Furthermore, he says that many insurers, like Mapfre itself, are just at the start of their journeys. "If I had to say, how far we are grabbing the AI opportunity from one to 10, I would say two".

He adds that "the opportunity is brutally, brutally high", but so are the organisational difficulties in capturing it.

Measuring AI impact and investment

Figures for how much insurers are investing in AI remain unevenly disclosed and cherry picked at best. Some insurers are explicit: US-based insurer Nationwide announced a $1.5bn technology and AI investment through 2028, while German insurer Ergo has committed €130m ($153m) to developing generative AI technologies.

Most, however, point instead to big ambitions and partnerships, like Allianz announcing a global collaboration with Anthropic to develop AI tools across claims, underwriting and employee workflows.

Caillat says Axa is now moving from experimentation towards focus and scale, while Qeli says Swiss Re has "doubled down on AI" over the past year.

Hard performance metrics remain rare and nebulous as well. AIG is one of the few large insurers to publish clear numbers, claiming its AI-powered underwriting platform AIG Assist has helped deliver a 30% increase in quoting submissions, a 55% reduction in time to quote and approximately 40% increase in binding of middle market property cover at its Lexington subsidiary.

Allianz has disclosed some operational figures, saying fully automated processing via AI accounted for 49.7% of German pet

AI progress at a glance

Swiss Re says AI is widely used and is allowing the reinsurer to process highly detailed policy or claims data files and other document-heavy workflows that were previously too costly to analyse manually.

Zurich says AI can accelerate underwriting and claims, while it is already being used to improve customer service and deliver better risk information. The insurer has also launched a Zurich AI Lab with ETH Zurich and the University of St Gallen to accelerate the transformation of the insurance business model.

Mapfre says it has around 130 AI use cases across its European, US and Latin American businesess, with most current value concentrated in fraud detection and pricing.

Axa says AI is used across document summarisation, underwriting, claims triage, customer journeys and software development. Axa has 3.5bn documents containing data, which it is "gradually using generative AI to dig into".

Markel says AI is embedded widely and is now "enhancing or empowering" much of its risk function's daily work.

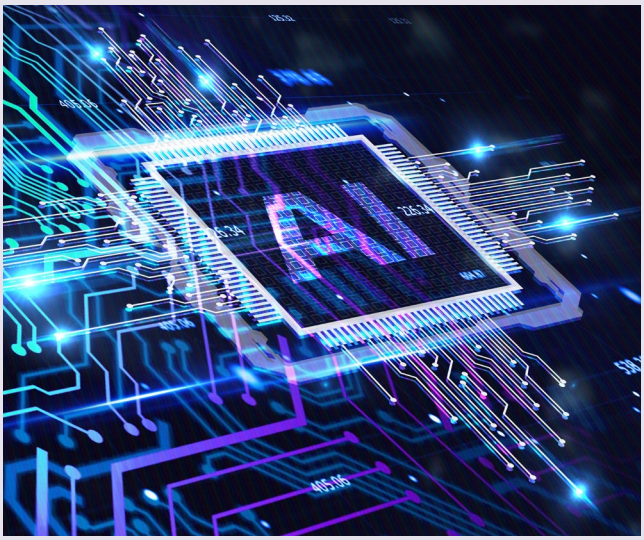

Munich Re says it will "reduce complexity and combine our market-leading know-how with artificial intelligence to boost our speed" under its 2030 strategy plan, which is targeting €600m in cost-savings.

Allianz says AI is being scaled across claims, underwriting and employee workflows through its global partnership with Anthropic. The insurer has also disclosed more than 900 registered AI use cases internally, including automated claims handling in German pet insurance.

insurance claims in 2025, with simple claims paid within a few hours.

Aviva, meanwhile, told investors its motor claims transformation saved more than £60m ($82m) in 2024, after cutting liability assessment time for complex cases by 23 days, improving claims routing accuracy by 30% and reducing customer complaints by 65%.

The human centre of underwriting

Some of the most advanced examples of AI use have been in underwriting. Peter Zaffino, AIG chairman and chief executive, says the American insurer started with the transformation "at the core of our business in underwriting, where we feel the impact will be most profound".

In Zaffino's more detailed account of his AI philosophy, he stresses that machines will not be wholly replacing people. "Our strategy in AI has been built on five pillars, but the middle is the underwriter," he says in an interview with *Time*. "And the underwriter is still making underwriting decisions and going to make risk and policy decisions."

Caillat says Axa takes a similar human-centric approach. "We like to have the human in the loop to make sure that things are under control, but it doesn't mean that you have to do everything from A to Z," he says.

The view from below

Leaders across the insurance industry have often heralded the many advantages of AI and described it as a revolutionary technology. How much of those benefits have managed to have an impact on the ground?

Speaking off the record to *InsuranceERM*, two employees at a major European reinsurer describe an AI push that, in their view, can appear driven as much by market messaging as by practical utility. "It feels like sometimes the leadership is trying to figure out how they can say AI as often as possible to please the shareholders," one says.

For now, they argue, material impacts remain uneven. Most use cases still centre on rewriting emails, summarising documents and speeding up coding, rather than reshaping the core actuarial or underwriting work of the business.

Another employee says many actuarial and underwriting tools currently offer only very rudimentary AI features, with much of day-to-day work "remaining as it used to be".

Across the industry, early-career insurance employees describe anxiety about what AI could mean for their jobs.

Every insurer interviewed for this feature stressed a human-in-the-loop approach.

However of the five early-career employees spoken to, all without exception expressed a fear of AI-driven redundancies at their respective companies and a doubt that corporate promises of human-centred approaches will hold against market pressures.

"The last thing I would want if a family member dies is calling my insurer just to talk to AI"

Alison Martin, Zurich

The same applies in life insurance, according to Alison Martin, Zurich's chief executive of life, health and bank distribution. She gives the example of AI tools helping an underwriter understand a complex medical history.

"You've got very complicated medical information about someone's heart condition," she says. AI can help "ingest all of that and summarise it".

She acknowledges that the boundary of where AI would be allowed to make decisions may move over time in lower-risk areas. Martin says Zurich already has examples in claims where "on certain small size claims in the non-life retail business, we've already tested letting the AI make the claims decisions".

Customer contact is the most visible test

If underwriting is where AI may improve margins, customer engagement is where it may be most visible. Insurers have announced significant AI pushes in customer support, where some roles are set to be replaced by the technology.

Caillat says the clearest change for policyholders will come through smoother journeys and more valuable interactions.

"We want all the admin parts to become smoother," he says. "The contacts we have with our clients [should] be more focused on what

"Chubb is embracing the power of technology, AI, data and business process change, transforming ourselves so we thrive for years to come"

Evan Greenberg, Chubb

brings them value. So it's advice, it's prevention, it's understanding them better."

Zurich's Martin accepts that AI has uses in customer service, but she is wary of inserting it into emotionally sensitive conversations.

"The last thing I would want if a family member dies is calling my insurer just to talk to AI," she says.

For Zurich, there is a distinction between simple administrative interactions and vulnerable customer moments.

She adds: "There is insufficient evidence that agentic AI in a call centre can interact with the level of empathy that we expect humans to be able to do."

The workforce question

Every AI strategy eventually reaches the workforce question, and who – if so much efficiency has to be achieved – will be laid off along the way.

Caillat says Axa sees AI as a way to redirect work towards customer value, not as a redundancy programme.

He says advice and prevention remain areas where Axa can expand. "Being able to redirect careers and jobs to make sure that we bring more value to our clients is a fantastic opportunity."

Qeli frames the issue differently for reinsurance, where expert judgement, negotiation and client relationships remain central.

"There will be a shift in tasks," he says. "Eventually for certain roles you might also have more radical shifts also in the tasks involved."

But he cautions against treating every restructuring effort as purely AI-driven. "It's a bit like the ROI [return on investment] question for AI," he says. "How much of that is attributable to AI as in contrast to plans that certain companies might have for restructuring anyway?"

Yet the job cuts have already begun. Over the past year, Allianz, Ergo, Munich Re and Chubb have either announced AI-linked redundancy plans or set out workforce reductions tied to AI.

Risk teams enter the AI debate

For chief risk officers (CROs), AI is seen as both a tool and a source of risk.

Henry Gardener, CRO at Markel, says: "AI is extraordinary and it is changing the nature of work for Markel, for risk teams and for our customers. We are focused on understanding, using and learning AI."

He adds that "many AI risks are new flavours of existing risk", but also warns: "It is changing the pace and scale at which decisions are made."

AI can make familiar weaknesses more dangerous: poor governance, weak documentation, unclear accountability, biased data and excessive reliance on tools.

Bart Frijns, NN Group's interim head of the enterprise and non-financial risk management department, says this risk is already visible. "If you have a tool that really works well, people get a bit lazy," he says. "You stop thinking."

For risk teams, he says, that cannot happen. "For risk management, this is one of the key capabilities we need," Frijns says. "Be critical towards AI."

Qeli questions how junior staff will handle tools which can produce seemingly infallible results at speed. "How do we preserve the critical thinking of our experts in the future if they have tools that are much more powerful than the tools that they had earlier?"

The risk is moving too slowly

The leaders interviewed speak about speed, relevance and control as central to their AI transformation plans.

But there are challenges in trying to move faster without losing judgement. Insurers want quicker and more accurate underwriting, claims and distribution – but not weaker accountability. They want efficiency gains – but not hollowed-out expertise.

However, of all the risks, losing competitiveness is perhaps the biggest of them all.

"AI adoption is happening everywhere in the insurance industry," says Frijns. "For me, it's a risk if we as a company are not [on top of it]."

In 2022, the tech industry did not stay out of ideas for long, and tech companies that didn't adapt at the time are struggling now. Standing still, it turns out, was not the correct move. Insurance may not have that luxury either.

AI in the insurance value chain

Insurers have integrated AI into many aspects of their operations. Anna Sagar presents a series of case studies showing how it has been done, and discusses the secrets to success

Artificial intelligence is the talk of the insurance town, with expectations that it will revolutionise the insurance value chain in the coming years. But where and how are insurers actually applying AI in their businesses?

One of the most recent surveys of the market, performed by the Swiss Re Institute, found of the 187 AI use cases it analysed, underwriting was the dominant area, closely followed by operations (see Figure 1).

In 2023, the European Insurance and Occupational Pensions Authority (Eiopa) polled 209 re/insurers from 22 countries and found customer service was the area of the insurance value chain where most insurers are using AI, followed by fraud detection and claims management.

This article describes five case studies where AI has been embedded in the insurance value chain, how they were successfully implemented, and what insurers looking to integrate AI should consider.

Underwriting – Generali

Italian insurer Generali teamed up with AI technology vendor Sixfold to improve the efficiency of Generali's cyber underwriting, enabling the firm to grow its premiums with the same underwriting team.

The Sixfold platform automates the more time-consuming parts of the underwriting process, such as reviewing large volumes of technical documentation, questionnaires and control descriptions, while remaining aligned with Generali's underwriting guidelines and strategy, as well as data privacy and auditability.

The firms ran an initial pilot which aimed to cut the administrative work performed by underwriters and accelerate responses to clients and brokers.

The initiative also aimed to improve cyber risk engineering activities (the process of identifying risks and then suggesting mitigating actions and controls), implement a structured and automated risk rating methodology, and build a data lake directly extracted from submission documentation to support data-driven portfolio and business steering.

The firms said the results were immediate and tangible. Risk reviews that could take hours or days were automated and allowed underwriters to come back to brokers much faster. Cyber risk engineering teams could product reports in a few hours, compared to the previous two-day manual process.

Now, most cyber submissions use the Sixfold solution, with turnaround times for distribution channels cut by 50%.

Generali's underwriters also found it allowed them to focus on building business relationships and making judgement calls, which is where underwriting expertise really adds value.

To ensure the project was implemented properly, underwriters from different markets were involved from the early stages of the

Figure 1: Global AI use cases by value chain (January 2024-October 2025)

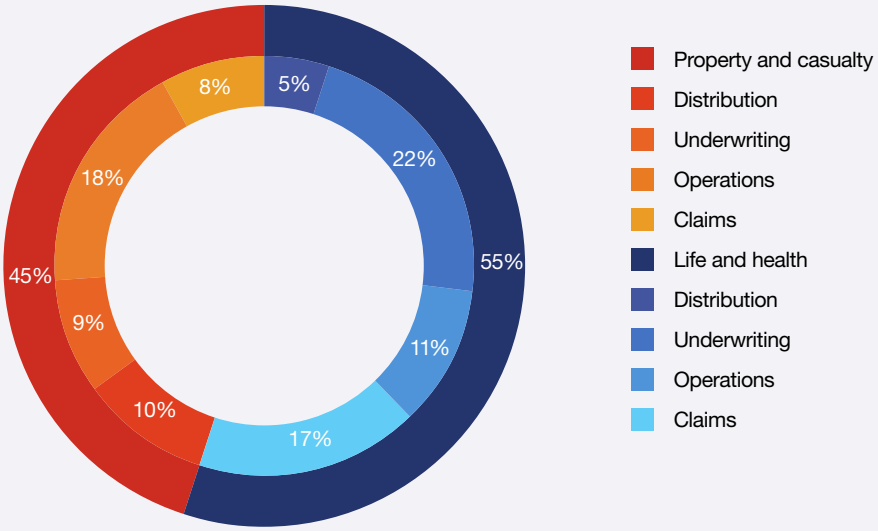

Source: Swiss Re Institute, sample of 187 use cases from major insurers in Apac, North America and Europe

pilot, with their feedback incorporated into the solution.

Generali said the biggest challenges around implementation were human and organisational, as introducing AI raised questions around trust, accountability and control.

The transparency and auditability features in Sixfold's software, as well as its strict approach to data sovereignty, helped address these challenges.

Underwriting – Hiscox

Phil Withey

Similar to Generali and Sixfold, insurer Hiscox teamed up with Google Cloud in 2024 to create an AI-enhanced lead underwriting model in the London insurance market

The project initially focused on automating the underwriting process, from submission to quote, in the sabotage and terrorism insurance markets.

The aim of the project was to help underwriters make faster, more confident decisions on complex risks, enabling Hiscox to serve its customers more effectively and efficiently.

It used Hiscox's in-house AI models and Google Cloud's AI platform to extract key data and insights from email submissions to automatically feed the underwriting and pricing process.

This was traditionally a manual process that would have taken up to three days. Hiscox says the technology reduced the time-to-quote to minutes.

Phil Withey, chief technology officer for Hiscox London Market, said leveraging AI can deliver "tangible benefits for customers, strengthen relationships with brokers, and free up our underwriters from manual tasks – allowing them to focus on more complex risks where human expertise is critical".

Hiscox has gone on to widen its collaboration with Google Cloud, using generative AI in its underwriting and risk modelling. Further product lines, beyond the initial sabotage and terrorism, are now using the system, according to Hiscox.

The firm said it adopted a "collaborative approach" to ensure effective implementation, with underwriters, technologists and governance teams working together from the outset.

There was also an emphasis on transparency, explainability and user trust, alongside appropriate controls and oversight.

"By involving people at every step of the journey, the project was able to balance technological advancement with human interaction and with risk management truly embedded in the thinking," Withey said.

Customer service – Allianz Partners

Insurer and assistance company Allianz Partners has been using agentic AI in its roadside assistance business since 2025.

The primary goal of using agentic AI was to improve the customer experience by cutting wait times and increasing service availability, especially during peak travel periods when call volumes can rise. The firm said productivity and efficiency gains were a secondary outcome.

The insurer uses multilingual voice agents across the whole assistance journey, from the initial call to dispatching the appropriate support.

The digital agent assesses context such as vehicle type, location, time of day and customer status to determine next best steps, with sensitive or complex cases automatically escalated to a human colleague.

Figure 2: European use of AI in different areas of the value chain (June 2023)

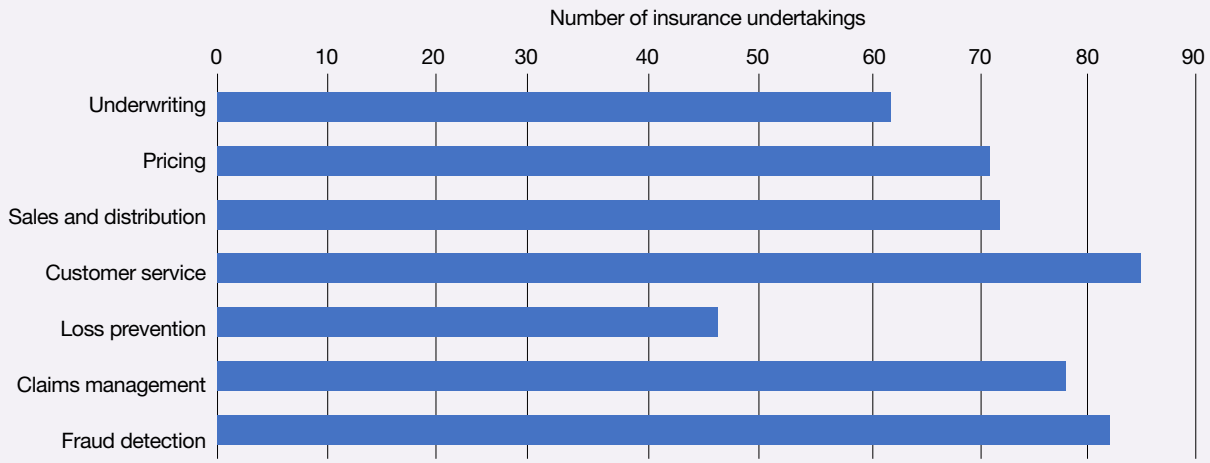

Source: Report on the digitalisation of the European insurance, Eiopa

Matt Crawford

Allianz Partners said using digital agents significantly improved peak-time call handling and ensured consistent service during busy periods. Across global markets, the firm claimed customer satisfaction scores for the AI-supported interactions average 4.5 out of 5, on par with human-only service.

Matt Crawford, head of UK roadside assistance sales at Allianz Partners, said: "Unlike traditional automation, agentic AI can handle routine cases end-to-end – verifying customer and coverage details, selecting an appropriate service, dispatching support and following up with the customer automatically."

"In roadside assistance, response times play an important role in shaping the overall customer experience. Agentic AI gives us a practical way to stay responsive during peak demand, while keeping human support readily available."

Allianz Partners said implementation was backed by a group-wide governance framework for responsible AI use. This includes a mandatory "human-in-the-loop" principle to ensure that final decisions in escalations, exceptions and sensitive cases always sit with human employees.

The insurer added that all its AI systems are "designed to be transparent and fully documented", and customers can request a human agent at any time.

"Responsible AI frameworks, including transparency and human oversight, make it possible to scale operations during high-demand periods without compromising service quality or trust," Crawford added.

Claims management – ManyPets

Pet insurance provider ManyPets has developed Millie, an AI-backed claims assistant, which it says is part of its long-term strategy to improve claims handling.

Pierre du Toit

The firm said that each year it receives hundreds of thousands of claims, so it leveraged AI to help deliver faster, more consistent outcomes for customers. It said Millie is "not a standalone project" but is at the centre of its claims operation.

ManyPets' chief data officer, Pierre du Toit, said: "We wanted to bring consistency to claims decisions, so customers know what to expect, and to reduce friction during treatment, when speed and clarity matter most. By combining automation with human expertise, we set out to remove unnecessary delays and give pet owners greater peace of mind."

Integrating Millie into the process reduced claims processing times, according to ManyPets, with most claims paid within one day, and 80% within five days. This is almost double the rate it achieved a year ago.

Currently, all claims go through automated assessment, and up to 55% are fully processed by AI.

Du Toit added that Millie is helping the company "move towards a real-time claims service".

Other benefits have included improvements in Trustpilot reviews as well as customer sentiment and retention.

To manage implementation, the firm said it built a multi-disciplinary team including product managers, data scientists, engineers, claims handlers and veterinary experts to redesign the claims process.

"Rather than automate existing manual workflows, we rethought them from the ground up for an AI-driven approach. We also invested in high-quality claims data, robust infrastructure

to deploy models, and strong monitoring to ensure consistent performance.

"This combination of process redesign, data investment and cross-functional collaboration was key to making the system work reliably at scale," du Toit said.

Fraud detection – Covéa UK

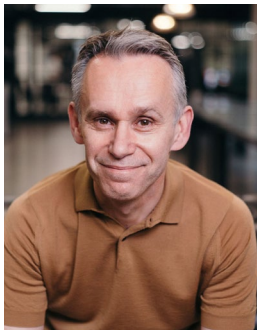

George Robbins

Home, motor and commercial insurer Covéa UK partnered with technology vendor Shift Technology to transform its approach to fraud detection, moving from "fragmented, rule-based processes to a unified, AI-driven risk intelligence framework".

The firm said it was facing increasing complexity in fraud and financial crime detection across its business lines.

Disconnected systems and tools led to inconsistent decision-making and late identification of fraud – often only after a claim was well underway or even already paid. Fraud controls were reactive rather than proactive, leading to inefficiencies, leakage and operational strain.

The firm deployed Shift's AI-powered platform to create a single view of risk across underwriting and claims.

The platform ingests data from internal systems and external sources like Companies House, and the AI model analyses patterns to generate "explainable risk signals".

Once it has ingested the data, the model will score and explain why it thinks a case is suspicious, allowing the claims handlers to prioritise potentially higher-risk claims for examination. Covéa UK said human decision-making remains central to its process, with AI looking to "support, not replace, claims handlers".

In order to achieve the cultural buy-in necessary to integrate the platform effectively, the firm adopted a "test and learn" approach, where claims handlers were involved in refining the model and ensuring the outputs aligned with their expertise. Covéa UK also established governance processes for ongoing optimisation.

All of this ensured the solution was embedded into day-to-day operations rather than operating as a standalone tool.

According to Covéa UK, the platform improved the hit rate – the proportion of claims flagged as needing further investigation – to 62%. The industry standard hit rates are around 20% to 30%.

The initiative ultimately led to £1.5m saved in policy repudiations in Q4 2025 alone. Return on investment was achieved in three months.

Shift Technology said the solution improved collaboration between underwriting and claims teams, enabling a more consistent and proactive approach to risk management.

George Robbins, head of UK markets at Shift Technology, said: "Covéa's approach illustrates a broader shift in the market, where insurers are no longer looking at fraud detection in isolation but as part of an end-to-end transformation agenda.

"By combining predictive models to surface risk, generative AI to synthesise and explain cases, and agentic capabilities to orchestrate actions, always with human oversight, insurers can take earlier, more consistent actions and generate measurable value at scale."

Implementation hurdles

Meghan Anzelc

From the above case studies, some key themes for effective implementation have emerged.

Having a "human in the loop" and cultural buy-in from employees is vital. Dropping in a solution and mandating its use will not set up a project for success.

Including employees in the creation and design means that it will address problems more effectively as it is more integrated into day-to-day work but it will give employees more agency.

As part of this, having an identified pain point or specific problem that you want to address with a specific type of AI – whether that is narrow, generative, agentic or machine learning – will ensure that AI projects are targeted and effective.

To bring all of the above together, having a governance framework around AI will mean that insurers can move safely from experimentation towards structured adoption.

The Lloyd's Market Association this year brought out an AI adoption toolkit, which centres on five core principles: governance and accountability; risk tiering to determine appropriate levels of control; data protection, security and intellectual property; training and awareness; and pragmatic adoption.

The key takeaway from the toolkit is that starting with lower-risk use cases, embedding controls from the start and scaling adoption gradually will ensure that AI adoption is managed gradually.

The case studies above also discuss how AI solutions should be "future-ready". The AI landscape is constantly evolving, so solutions need to have the ability to advance with the technology.

Reviewing and evaluating outputs, as well as looking at what developments are going on in the wider landscape, will ensure that progress continues at pace, rather than in a stop-start manner.

Meghan Anzelc, global leader of transformation solutions at Aon's Strategy and Technology Group, said that most insurers claim to have deployed AI but fewer "can point to results".

She said less than 15% report measurable impact on combined ratio, cycle time or loss ratio.

"That gap isn't a technology problem. It's an operating model and change problem," she said.

Anzelc pointed to a carrier Aon had worked with where pricing and underwriting were rebuilt around an AI-enabled pricing engine, which was "technically excellent" but faltered due to adoption issues.

"Nothing failed. The model worked. What failed was adoption. Underwriters were asked to trust outputs they hadn't helped shape, with limited clarity on how the model should influence judgment, authority or accountability," she explained.

Anzelc said that defining the problem more clearly, especially what "good enough" meant, would have saved effort.

"Insurers that scale AI treat adoption as an operational requirement, not an afterthought. Without ownership, incentives and clear expectations for use, even the best models become unused – or quietly ignored," she said.

Harnessing the power of AI for actuaries

AI offers actuaries powerful opportunities to improve efficiency and insight. But as technology advances, how do you maintain the control over data, systems and processes that's so critical to decision-making and compliance? Martin Sarjeant, Head of Solutions Management for Insurance and Climate Risk at FIS, explores ways to embrace AI without sacrificing governance, resilience and confidence in results.

How is AI transforming the world of actuarial modelling?

The recent progress of AI and the proliferation of ever changing open-source code, in theory means you can build example actuarial code quickly. With large language models, you can accelerate coding tasks, improve predictive analytics and process unstructured data, simply by writing prompts in natural language.

Martin Sarjeant

As inflation persists globally and operational costs continue to rise, insurers are looking at AI to drive cost-savings. By both automating and streamlining complex workflows from end to end, AI tools can help actuaries step up productivity and ultimately do more with less.

What are the risks of relying on open-source code for modelling?

AI may be advancing quickly, but it doesn't always have your back. For calculating premiums, solvency, complying with regulation and protecting stakeholder interests actuaries can't afford to make errors so when you're building and running models, you need to be sure that your code is validated, and your calculations are consistently on point.

Whilst generative AI can help you write example code without much coding expertise, there's a large investment required to take example code and build models and a platform for production use. The overall solution needs to be maintained, trusted, versioned, correct and be developed without use other's intellectual property.

It's also important that your models are easy to adapt and scale to new products and future reporting requirements. Plus, your actuarial operations as a whole must be resilient. That means knowing that your systems are available and supported around the clock, your processes are rigorously governed, and your data is safe from the growing threat of cyberattacks.

Why is an established technology platform still the best choice for actuaries?

With a proven actuarial modelling and risk management solution such as FIS® Insurance Risk Suite – Prophet, you get much more than modelling software. You also gain access to the expertise, strength and scale of one of the largest financial technology companies in the world.

All of our library code is validated by actuarial experts, trusted by the industry, so you can run reliable risk models out of the box, and easily adapt them to your specific needs. And with global support, you can count on our 600+ strong team to help you meet your latest business and regulatory challenges.

You can also rely on high levels of availability, best-practice cybersecurity standards and full compliance with your firm's AI and IP policies. Whether you access the solution in a secure pay-as-you-go SaaS setup or as a managed service, you can easily scale Insurance Risk Suite – Prophet up and down to meet changing regulatory and internal risk management needs.

Where is FIS applying AI technologies to help improve productivity and reduce costs?

At FIS, we're introducing a range of AI productivity tools and workflows to Insurance Risk Suite – Prophet, designed to help our clients drive automation, control and savings across their actuarial modelling operations.

We've already added a chatbot-style AI assistant that will give you instant answers to any question about using our solution. And we're helping transform the way that actuarial teams capture, structure and communicate knowledge with an AI solution that can turn complex code, procedures and collective expertise into clear, professional and audit-ready documentation.

FIS is also investing in building an AI coding assistant that will help you create model code faster, amend models, reduce errors and optimize your models. Our goal is to allow you to generate code and build models using natural language instead of code syntax. In the meantime, we're supporting the wider application of AI to setting up model runs and recommending the right machines for each run.

Our natural language vision for the Insurance Risk Suite extends to all areas, including process and data. So, users can ask questions of the data in natural language rather than needing to write complex queries, validate assumptions, create test models and so on, we are also embracing agentic AI throughout the end-to-end actuarial processes.

https://www.fisglobal.com/products/fis-insurance-risk-suite

What insurers should consider when deploying AI

Pete Carvill investigates the key decisions firms must make when embedding AI tools in their business

There has been a swift and widespread move among insurers to embrace AI. In 2024, Eiopa's comprehensive survey[1] of European insurers found that 50% of non-life insurers and 24% of life insurers were using AI for pricing and underwriting, fraud detection and claims management – percentages that have no doubt risen since then. However, there is still much to learn on the processes of how to embed this technology into the workings of an organisation.

As McKinsey & Co. wrote in its *The Future of AI in the Insurance Industry report*[2], "While AI holds immense potential for insurers, scaling it enterprise-wide remains challenging. Security risks, high costs, the risk of getting locked in with suppliers, talent shortages, cultural resistance, governance gaps and legacy infrastructure often hinder progress. A true transformation requires addressing these barriers head-on – and doing so in a thoughtful way that avoids creating 'tomorrow's legacy' with the proliferation of approaches and solutions we are seeing today."

There are differing viewpoints on the best approach, notably given by consultancies CGI[3] and the Boston Consulting Group[4]. Essentially, both approaches can be boiled down to a handful of principles: knowing what parts of your business you want to buttress with AI; testing and learning in targeted areas; focusing on high-impact opportunities; and ensuring support from both leadership and the people involved in executing this transition.

But there are few guidelines for insurers on how best to implement AI.

"From a standards perspective," says Rachel Reid, global co-head of AI and co-lead of global cybersecurity and data privacy at law firm Eversheds Sutherland, "there's no single, binding framework that answers every question for insurers today, but there are increasingly clear signposts. Leading insurers are drawing on a combination of sources: emerging AI governance standards (such as ISO/IEC AI management standards), sector-agnostic risk frameworks (such as NIST AI Risk Management Framework), existing model risk management disciplines, and privacy and data protection laws."

To develop internally or externally

Near the outset of an AI project, there is a point when a decision is made about whether to look to develop AI capabilities internally or externally. Does an insurer look to build their offering from the ground up, making sure that each step and process is tailored to their specific needs? Or do they look for an external vendor with more experience, but who may struggle to tailor the product to the specificity required?

Rachel Reid, Eversheds Sutherland

Neil Covington, FIS

"This," says Neil Covington, director of solutions management at actuarial technology firm FIS, "is a decision that in our experience is one that's being driven by risk, differentiation and scalability. From what we've seen in dealing with our clients, firms need to build internally when the AI is tied to their proprietary data or internal processes that can't be standardised or exposed. The time to go external is when the use case needs enterprise-grade governance, security, auditability and long-term maintenance."

He adds: "These are the things that are costly to build and maintain if you're trying to go it alone. Our perspective, from our years in implementing these programmes, is that many insurers explicitly favour external, domain-specific platforms for regulated workflows while opting for internal experimentation for less-critical use cases."

Timing is also an issue, says Reid. She says her experience shows that insurers get into difficulties when they implement the decision at too early a stage.

"The most-successful programmes," she adds, "start with a clearly defined use case and risk tolerance, pilot with limited scope, and then decide whether to scale internally or source externally based on performance, control requirements and cost of ownership over time. The question is less 'build vs. buy' in the abstract, and more about asking what level of control, transparency and accountability this use case requires."

There is also a third route, says Manu Mazumdar, head of

Manu Mazumdar, Conning's Insurance Research

data analytics and insurance technology at Conning's Insurance Research: a hybrid option in which insurers use third-party AI platforms for foundational capabilities while layering on proprietary models, rules and controls.

He adds: "It comes down to whether AI is core to competitive advantage or a shared utility. Insurers typically build AI internally when it affects proprietary underwriting logic, pricing or claims decisioning. They're more likely to buy when the use case is horizontal; like document ingestion, call summarisation, etc."

How to assess vendors

Should a firm look to bring in an external vendor, there are still hurdles to overcome. A rapidly evolving market attracts many new players who arrive with bold marketing plans and claims on success. But these things remain hard to measure.

Instead, Reid says that insurers should look at these firms through the prism of insurance-grade risk management, asking practical questions such as: Has this firm deployed AI in regulated environments? Can it explain how models are trained, tested, monitored and retrained? Does it have defensible positions on bias, explainability, data provenance and human oversight?

"Experience," she explains, "is not just about how long a firm has been 'doing' AI, but whether it understands the realities of insurance operations: legacy data, regulatory exams, claims disputes, consumer protection obligations and enterprise governance."

Reid advises that markers of success could include concrete insurance use cases, mature documentation practices, willingness to support audits, and contractual terms that clearly allocate risk, liability and data rights.

She concludes: "In short, insurers should assess AI vendors the same way they assess critical service providers, but with heightened attention to transparency and ongoing monitoring."

The data

When deploying AI, and in particular large language models (LLMs), an insurance firm should look at what data it holds and assess what work needs to be done in order to 'cleanse' it so that it may be used.

"One of the constraints on AI adoption in the insurance industry," says Mazumdar, "is not the models but the data and the architecture that surrounds them. Many core policy, claims and billing systems were originally not designed to support real-time data access, unstructured data, or continuous model feedback loops. One of the bottlenecks to AI in insurance isn't the technology, it's legacy data and systems that weren't designed for AI."

Covington says that data quality and integration are the hardest problems, far surpassing those presented by algorithms, when it comes to AI processes.

"The challenges and responses we see that recur all the time," he says, "are things like integrating the new systems with the legacy systems. The reason for that is older platforms require Application Programming Interfaces (APIs), staged integration or parallel workflows to avoid destabilising core systems. You also have issues with data quality and ownership, preparing that data for LLMs, and making sure that all the AI outputs align with existing models, assumptions and controls. That's a tough thing to do; AI is expected to fit into existing governance, not bypass it."

The people problems

Tackling these problems means ensuring that the right people are in place. The consulting firm McKinsey & Co. recommends that 70-80% of digital talent should be in-house, although finding those people is another challenge in a highly competitive jobs market. But in waiting and holding out to see where the market is heading while trying to find 'ideal' candidates, insurers put themselves at risks in other ways.

"Because of uncertainty about how to proceed," writes the Boston Consulting Group, "many insurers are adopting a late-follower strategy. This is a mistake. Building the necessary capabilities – particularly the human skills – takes time, so slower movers risk falling behind speedier competitors. Companies that move more quickly will likely reap significant strategic benefits in the longer run as well."

All together now

In a 2025 paper[5], professional services firm Deloitte examined all the reasons that had made implementations of AI unsuccessful.

Tellingly, the firm's authors wrote: "[Among] respondents, the lack of business line support was deemed the most important factor that made implementation unsuccessful."

Conversely, successful implementations showed what the firm called a 'mirrored picture'.

Along with good data, Deloitte wrote: "[Close] collaboration across business, tech, data and talent functions was cited as the biggest reason for success among those surveyed. When these functions work in harmony, they can create a favourable environment for AI initiatives to succeed."

It would seem, then, that what is most needed for successful deployment before anything else may be the right attitudes from the top of, and throughout, the implementing firm.

Footnotes

1 Report on the digitalisation of the European insurance sector, European Insurance and Occupational Pensions Authority
2 The future of AI in the insurance industry, McKinsey & Co
3 The intelligent insurer: moving beyond legacy constraints with AI-driven modernization, CGI
4 How insurers can supercharge their strategy with AI, Boston Consulting Group
5 Are insurers truly ready to scale gen AI?, Deloitte

AI's incremental rise in actuarial modelling

The AI revolution is impacting actuarial modelling, but not as a big bang, rather an evolution as insurers adopt AI gradually and tactically. Ronan McCaughey explains

If there is a defining characteristic of AI's impact on actuarial modelling today, it is restraint.

Speaking to *InsuranceERM*, Neil Covington, vice president, product management executive for insurance at FIS, describes a market where the insurance technology firm's clients are adopting AI incrementally and tactically, with a strong emphasis on governance, auditability, and risk control – rather than full automation.

"Today, AI is being used primarily as a coding accelerator, not an autonomous developer," he says.

Covington paints a picture of AI-driven evolution over the next two to three years, rather than disruption of the whole actuarial modelling process. He says: "The trajectory is clear, but bounded by regulation and risk appetite."

At a high level, he explains that insurers are expected to move from AI-assisted tasks to AI-assisted workflows. This will mean shifting from isolated AI tools to AI becoming embedded in actuarial platforms.

Another change will involve moving from "help me write code" to "help me implement a governed change end to end". Amid this progress, Covington emphasises that several structural constraints will remain, and human sign-off on models will remain mandatory.

Assisted coding

Assisted coding with a human-in-the-loop offers a major opportunity for AI to help actuarial modellers.

Covington says the main areas FIS is initially looking to support include translating natural language specifications into syntax for the firm's actuarial modelling system; supporting incremental model changes by adding variables; and enforcing standardised coding patterns.

What insurers are not doing today, and what FIS is not initially looking to support, is allowing AI to generate complete production models without review and using public or external copilots directly against regulated model codebases, says Covington.

Documentation and model governance

Documentation and model governance are emerging as AI's strongest early wins, says FIS's Covington. He notes the main areas that FIS is initially looking to support include automated conversion of Prophet model artefacts and code into regulator-ready documentation.

Covington can also see a role for AI to maintain documentation in step with model changes, which would reduce manual effort and operational risk.

Neil Covington, FIS

Adrian Parris, MetLife

AI and coding at MetLife

MetLife shows how this evolution of AI in coding for actuarial modelling is unfolding in practice. Adrian Parris, vice president for US actuarial modelling at the insurer, says AI-powered coding assistants are helping MetLife to generate, refine and troubleshoot actuarial code. That reduces development time and faster iteration.

"These tools are enhancing the productivity of our associates, enabling us to be more efficient," Parris says.

He says the insurer is also leveraging AI for proxy modelling, streamlining processes and developing automated data and model management.

Rather than a big bang implementation, MetLife is adopting a careful approach: parallel runs, pilot projects and controlled experimentation.

Parris frames the value not as replacing actuaries, but as refocusing their effort. "This shift enables modellers to concentrate on testing, analysing results, and engaging stakeholders – areas where actuarial judgment adds the most value," he explains.

Overall, Parris says: "AI is transforming actuarial work at MetLife. In combination with top-down initiatives, our efforts are driven by grassroots experimentation. Associates are identifying practical use cases, fostering a culture of innovation and continuous improvement."

Open-source platforms

The development of AI comes amid a growing synergy between

the technology and open-source programming languages like Python in actuarial modelling.

Given that many modellers are coding AI tools in Python, alongside actuarial platforms, Covington says FIS expects deeper integration over the next two to three years, but as a complement to core actuarial systems, not a replacement.

Across FIS's insurer clients, he explains that Python is already used in controlled, adjacent roles, such as data preparation, cleansing and reconciliation, as well as exploratory analysis.

"The key point is Python is trusted for analysis and experimentation, but not as the system of record for regulated actuarial models," says the FIS expert.

AI may also help migrate tools and models built in out-of-favour languages, such as APL, Cobol and Fortran, into an open-source code set that opens up the possibility of more flexibility in how these tools are developed.

Agentic AI and the future

Looking ahead, there is genuine excitement in the insurance industry about the potential applications for agentic AI in actuarial modelling. Agentic AI refers to specialised AI agents that take responsibility for discrete, bounded tasks, such as querying, impact attribution, validation, reporting, and operating under explicit governance and human oversight.

In the first place, there is a view in the market that agentic AI will help bridge the communications and skills gap between modellers and non-specialists, and thereby make actuarial models more accessible to insurers' teams.

The Institute and Faculty of Actuaries' GenAI working party, has said the emergence of AI agents will fundamentally reshape how actuarial models are built and maintained. This point was made in a paper[1] published by Daniel Ramsay, co-chair of the working party and an associate director at consultancy Deloitte.

The paper argues AI agents will be capable of performing various modelling tasks and are likely to do so with greater reliability and quality.

"This is because AI can be instructed to ensure that the actuarial model artefacts they generate are always maintainable and fully comply with necessary best practices and regulatory standards. AI may achieve a level of consistency and precision that surpasses human actuaries, who, constrained by time pressures or inattention, may sometimes be forced to or inadvertently cut corners."

While the promise of AI agents in actuarial modelling is significant, the IFoA working party's paper stressed the need to recognise their inherent limitations, and the challenges involved in their practical deployment.

FIS's Covington says validation is where agentic AI is likely to be most transformative and most acceptable to regulators, providing agent-led validation with human sign-off.

This is because agentic AI enables specialised validation agents that can compare model versions automatically, detect anomalous movements in results and attribute changes back to data, assumptions and code changes.

Covington says the cumulative effect of agentic AI will lead to a different operating model for insurers, as per table 1.

Table 1: Possible impacts of agentic AI on actuarial modelling

Today	5-10 years
Actuaries write and debug	Actuaries design, supervise, approve
Validation is periodic	Validation is continuous and automated
Documentation is manual	Documentation is generated and maintained
Models are static artefacts	Models are governed, live systems

Source: FIS

Overall, it is clear that AI will evolve how models are built, but not who is responsible for them. The direction of travel points to the future of actuarial modelling not being defined by machines replacing actuaries, but rather actuaries orchestrating machines.

Agentic AI will take on more of the mechanical workload – coding, testing, documenting, validation – while humans focus on judgement, interpretation and decision-making.

The winners will not be those who automate fastest, but those who integrate AI most intelligently into governed, auditable workflows.

"Even over five to 10 years, several things are very unlikely to move," says Covington. These include full autonomy in regulated actuarial models; AI making unapproved assumption changes; black-box models replacing explainable ones and removal of actuarial accountability.

"This is explicitly reinforced in insurer discussions: trust is anchored in auditability, lineage, and human oversight, not AI capability alone," he says.

Footnote
1 Actuarial Modelling in the Age of AI Agents

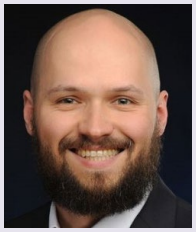

Alexey Mashechkin, head of technical excellence and analytics solutions, at Allianz Partners, who is also chair of lifelong learning at the Institute and Faculty of Actuaries' AI and Emerging Technologies Practice Board, stresses that while AI is a hot topic in actuarial modelling, its effectiveness also depends on insurers having good quality data.

He says: "This is the struggle. It is a struggle which can be supported with AI, but it cannot be solved and it's still a case of if you put garbage in, garbage comes out. This should be addressed by setting up proper data flows and governance."

Asked how black-box decisions made by AI and machine learning models can be made more explainable, he says there are analytical approaches, with some libraries available for this purpose, but [explainability] "is not functioning well for each and every algorithm and this is a challenge for actuaries given regulatory requirements".

Building confidence in AI-driven decisions

The complexity of AI models makes them harder to understand and explain than traditional actuarial models. But stakeholders continue to demand model transparency, so the sector has responded with new frameworks, tools and approaches. Christopher Cundy reports

Actuarial and risk teams are no strangers to requirements for managing model risk and validating models. But the growing use of machine learning (ML) and AI models, including large language models (LLMs), is posing fresh challenges to conventional approaches used to gain confidence in modelling.

AI models typically have more autonomy than traditional models and are updated more frequently. The mathematical basis for AI models is far more complex, making them harder to explain to stakeholders. This view of AI models as a "black box" is doubly challenging when there are heightened concerns around data ethics, bias and fairness.

Insurance regulators, standard setters and associations have recognised these issues and produced rules and guidance around the use of AI systems and establishing AI governance frameworks (see box).

The Lloyd's Market Association (LMA) is among those bodies releasing AI guidance. As part of producing its AI Adoption Toolkit, it interviewed 11 Lloyd's managing agents, revealing: "Across all the interviews, one theme stood out clearly: there is no consensus around what an ideal AI framework would look like."

This reflects both the infancy of AI governance and how frameworks need to be proportional to the business and the models in use. The question of who oversees AI operations is a case in point. The LMA found some firms where the chief technology officer or chief operating officer has prime responsibility; some where AI oversight is allocated within the existing structure (e.g., the risk committee); while others had a dedicated AI governance committee.

Matthew Byrne, NFU Mutual

There are clearly divergent approaches, but it is worth stressing that the ever-expanding use of AI models within insurance businesses means oversight will become impossible for one person to manage, and will become the responsibility of everyone.

"Today, governance is no longer a single person monitoring a model. AI governance is closer to the concept of cyber security than the concept of simple monitoring of a model," said Claudio Senatore, an actuary and vice-chair of the Actuarial Association of Europe (AAE) AI and data science working group, at an AAE webinar in April.

Explaining models

Being able to explain how a model works and why it produces its outputs is key to gaining confidence in the model and using it to make business decisions. A huge amount of work is therefore going into making AI models transparent and showing they are accurate, fair and free from bias.

Selection of rules and guidance on AI from regulators, associations and standard setters

NAIC Model Bulletin: Use of artificial intelligence systems by insurers
Eiopa: Opinion on AI governance and risk management
ABI: Practical ideas for getting started with responsible AI
NIST: Artificial Intelligence Risk Management Framework (AI RMF 1.0)
Lloyd's Market Association: AI Adoption Toolkit

Model explainability methods

"Global" model-agnostic methods, focusing on general trends:
- Partial dependence plot (PDP)
- Accumulated local effect plots (ALE)
- Functional decomposition
- Permutation feature importance

"Local" model-agnostic methods, focusing on interpreting individual predictions:
- Individual conditional expectation curves (ICE)
- Local surrogate models (LIME)
- Kernal Shapley Additive exPlanations (kernal SHAP)

Ironically, the primary advice from the actuarial modelling community is to continue using interpretable models (such as the generalised linear models commonly used in pricing) if the data is structured and their accuracy is sufficient.

But since AI models can offer improved accuracy, and are better suited to using unstructured data, it remains necessary to make them interpretable. Data scientists have come up with several techniques (see box) that can be applied, but none are a silver bullet.

"The results let you peer into the models, so you can see one interpretation of how the model has arrived at that decision," says Matthew Byrne, head of model risk at UK insurer NFU Mutual, and a member of the AI and Emerging Technologies Practice Board at the Institute and Faculty of Actuaries.

"It is important for practitioners to be aware that these techniques definitely help, and it's better to do this than nothing, but each of the techniques have their own limitations. None are perfect – and that's probably why there are so many."

In the case of LLMs, there are a host of observability tools, such as Braintrust or LangChain, that can track prompts and responses and measure the quality of outputs. These become particularly important when using agentic AI tools.

"Observability is when you want to see what's actually happening at each stage of a long LLM or agent-driven process," explains Ronald Richman, founder and CEO at insureAI, a provider of AI-based actuarial tools, and an experienced actuary and risk manager.

"These platforms are really nice for giving you that insight into what's happening. It takes it away from being a black box and makes it very observable."

At a deeper level, there is the field of mechanistic interpretability, that tries to dissect what happens with the LLM, to make it less like a black box.

"Mathematically, an LLM is a huge set of parameter matrices that are multiplied one after the other. Mechanistic interpretability researchers try to find structure within those parameter matrices that, for example, are responsible for the friendliness of the LLM," says Richman.

The importance of this approach has been highlighted by various LLM phenomena, such as the commonly discussed hallucinations. Just as concerning is "alignment faking", where the LLMs understand they are being observed or audited, and give deceptive results to make it look like they are compliant with their training.

Richman also notes the drive towards designing AI models that are inherently interpretable – which has been nicknamed as "glass box" models.

"If you can find a model that's almost as strong as a black-box model, but it gives you full insight into what the model is actually doing, that's extremely powerful because you can change what the model is doing and impose your actuarial judgement," he says.

Human in the loop

One of the common approaches to assuring confidence in AI models is to include a "human in the loop". This implies some kind of manual intervention, either at predefined points in a process to decide whether the AI results should be implemented, or when it has to be escalated to a human because the AI cannot cope.

But AI systems are typically complex, autonomous and fast, so including a human in the loop can become practically very difficult.

Byrne explains: "In the case of typical actuarial models, for something like personal lines pricing, it's clearly impractical for a single human or even a team of humans to manually check every single price that comes out of that model before you show it to the customer. So, firms need a good suite of testing in place to validate that the model gives robust and reasonable results, and checking that the results are explainable."

Some of this is standard model risk management practice, but for AI models, Byrne says additional monitoring is needed to ensure they are fair and unbiased, and do not degrade or drift over time.

Richman says his philosophy when developing insureAI's reserving tool was to build a 'recommendation engine', where the actuary is given information to judge which reserving model is most appropriate, but also be able to tweak the model parameters.

"It's about being able to expose what is the decision-making basis of the model, and on that basis, the user can then interrogate it."

Communication

Communicating explainability and model risk management to stakeholders is the final challenge for AI models. For the experts contacted in the article, the approach tends to be bringing it back as close as possible to what boards, regulators and customers are familiar with.

"For our pricing work, for example, we use all sorts of different black-box models. But as long as you can re-express that in a manner that satisfies your stakeholders and reflects the underlying decision-making basis of the model, well, then you're in a great place," says Richman.

AI and ML models can offer distinct advantages over traditional models. But they are complex, techniques are evolving fast, and regulators are putting them under greater scrutiny. Insurers will need to constantly scrutinise the explainability and fairness of models, and raise the bar on model risk management.

As Byrne sums up: "With AI models it's about finding a sweet spot when the risks don't outweigh the benefits."

The role of CROs in managing AI risk

AI can bring significant advantages to insurers, but also new risks. Martin Assmann speaks to chief risk officers and AI experts to understand how they can be managed

The growing use of AI within insurers is forcing risk teams to rethink how the business is exposed and who, at the end of the day, owns decisions and mistakes made with machine assistance.

Asked who owns AI risk, ChatGPT offers a neat, but uninspired answer: "The business function deploying the AI owns the risk, with oversight from risk, compliance and the board."

However, *InsuranceERM* did not want to take the risk of relying on it. So we also asked leaders across the industry how ownership of AI risk is being shaped in practice.

AI risk is not one risk

AI risks are as multi-faceted as the answers a chatbot can give to the same question. They can encompass model, conduct, operational, cyber and data protection risks, for example.

Henry Gardener, chief risk officer (CRO) at Markel Insurance, tells *InsuranceERM* AI risks fundamentally "are human risks, either sped up or just made more systematic".

He explains this makes managing AI risk less complicated than it first appears, but also more difficult to contain.

"What AI changes is how quickly those risks can emerge and how widely they can spread if they are not well understood and managed," he adds.

Bart Frijns, NN Group's interim head of the enterprise and non-financial risk management department, reaches a similar conclusion from inside the risk function. He says the traditional CRO model, centred on frameworks, controls and periodic review, is no longer enough.

"You had a control framework, and you had to ensure these controls were executed well," Frijns says. "That's not the world we live in anymore."

Owning the decision

ChatGPT's first comment on AI risk looks correct: the first line usually owns the business process and that has not changed with AI.

Gardener says Markel's approach to oversight starts with the question of "who is the decision maker, and it's got to be a human decision maker".

He says responsibility for AI strategy sits close to the first line. "We would expect a divisional president to be taking responsibility for AI in their division," he says.

But Gardener says AI strategy, controls and guidance still need to be "embedded in each of the divisions", with risk "supporting and challenging" where needed.

> "What AI changes is how quickly those risks can emerge and how widely they can spread if they are not well understood and managed,"
>
> **Henry Gardener, Markel Insurance**

At Swiss Re, Ermir Qeli, head data science and AI, group data services, says AI governance is directly built into the way projects are framed.

"The guidelines are not a check that is somewhere out there and a last step," he says. "The second line and the first line collaborate very strongly when the guidelines are put together."

So the second line does not need to become the first line, but it also has to adapt.

Speaking at the Insurtech Insights conference in London in March, Camilla Bennett, CRO at Axa UK, said risk functions already have one structural advantage: they are used to looking across the whole organisation.

"The risk function or the CRO is the only other role, besides the CEO, that actually has always had to think about everything across the organisation," she said.

However, she added that "if we're going to be giving an opinion on technology, we actually need to understand it ourselves".

David Smith, CRO at Ardonagh Speciality, said at the same conference that the second line should not try to replicate every

"You had a control framework, and you had to ensure these controls were executed well. That's not the world we live in anymore"

Bart Frijns, NN Group

specialist capability already sitting in the business.

"It's not a good use of resources to duplicate subject matter experts in the second line with first line," he said. Instead, risk teams need enough technical understanding to challenge the first line and identify gaps.

But the second line still needs enough expertise to challenge. "You do need one or two people in the second line who can ask the obvious questions and make sure you're highlighting the right gaps."

The risk of lazy judgement

The most immediate AI risk may be something quite simple: people believing the machine too readily.

Swiss Re's Qeli calls this the critical thinking problem. AI can give junior staff capabilities that were impossible only a few years ago, but that creates a training challenge.

Gardener says within the risk function, AI can be useful for challenge, scenario thinking and identifying blind spots, but only if humans remain responsible for the conclusion.

He gives the example of using AI to test an argument. "Where are my gaps?" he says. "Why is this wrong? What have I missed?"

In insurance, junior underwriters, claims handlers, actuaries and risk analysts often learn through manual work: reading submissions, comparing clauses, checking files, reconciling data, writing reports and watching senior colleagues challenge assumptions.

Qeli says the answer is not to block junior staff from tools, but instead to "train and upskill people with the technology," along the way.

The upside

Most CROs now acknowledge that AI is already changing what risk teams can do and the profound effects it can have on the business.

Gardener cites an example where a Markel team was "able to grow their premium in nine months from $19m up to $50m, using the power of augmented AI".

He describes the model as "underwriter at the helm", rather than simply human in the loop. "You bring in human knowledge, human context."

AI then brings "additional analytics" and "additional inference", but every product still needs human judgement and accountability.

This echoes much of what *InsuranceERM* heard across the industry: that while AI changes a lot, responsibility remains shared. First line must own the decision, technology must own the tool, risk must own the challenge, compliance must own the regulation and senior management must own the appetite and strategy.

But no one can own AI risk alone.

Will CROs disappear?

What does being a CRO in the AI era look like? And if AI keeps improving, how long until the CRO is replaced by a risk adviser agent?

Markel's Gardener says the role is shifting towards improving decisions under uncertainty. "The pace of strategic decision making is increasing and often the degree of uncertainty with which those decisions need to be made is increasing as well," he says.

Markel has built a risk intelligence unit to help fill out what Gardener calls the "decision landscape" for boards, CEOs and business leaders. AI can help expand that landscape, he says, by testing arguments, finding blind spots and bringing in examples from outside insurance.

He points to studies suggesting AI can challenge boards effectively in controlled settings, while remaining cautious about turning that into a prediction that CROs will disappear.

"I still think you need somebody, a human in the centre of that to take that information, to take a decision on it, to own it and to turn that into truly constructive challenge."

For Gardener, the CRO's value lies not only in analysis, but in judgement, context and independent challenge. "I'd struggle to see how if you had an LLM or an agent that's just consuming the same information as the chief financial officer, and then has some prompts saying challenge this, that actually gets you to a better outcome," he says.

Gardener frames the CRO's AI role around confidence rather than constraint. "Our role as CROs is to help organisations move faster with confidence and build long-term resilience, not just adopt new technology," he says.

Rob White, chief risk and compliance officer at Acrisure, said during Insurtech Insights that the job is also becoming more human, requiring diplomacy, problem-solving and collaboration skills.

"The days of a CRO in particular that is an out-and-out technical specialist, that is consulted by the business on technical questions, but beyond that is probably a slightly more withdrawn, back-office function, I think are long gone," he says.

In collaboration with

InsuranceERM

FIS