

AAE Note

CLOSING THE AI GOVERNANCE GAP IN INSURANCE: THE ROLE OF ACTUARIES

AN AAE NOTE FOR INSURANCE EXECUTIVES

NOVEMBER 2026

AAE | ACTUARIAL ASSOCIATION OF EUROPE
ADVISING · ACHIEVING · ENGAGING

Closing the AI governance Gap in Insurance: The Role of Actuaries

When models become cheap, judgement becomes valuable

An AAE Note for Insurance Executives

AAE Artificial Intelligence and Data Science Working Group

Executive summary

Insurers are adopting AI faster than they can govern it. Regulators have turned that gap into a compliance problem, and recent, well-publicised AI incidents have turned it into a boardroom one too. AI does not reduce the need for actuarial expertise. It changes what that expertise is for. As models, code and sophisticated analysis become cheaper and more widely available, value shifts towards the professional judgement that runs through the whole modelling process, together with model-risk expertise and clear accountability. Closing the governance gap does not require building something new. Actuaries already run the closest equivalent under Solvency II, engage directly with the bodies shaping AI rules, and are increasingly building with the technology themselves. This note sets out that evidence, the pressure it is creating, and what insurers' leadership may wish to consider.

1. Where the market stands

AI adoption has moved from pilot to mainstream: EIOPA's most recent market-wide survey, published in February 2026 and covering 347 undertakings across 25 EU/EEA countries, found that nearly two-thirds of insurers already actively use generative AI. Most activity is concentrated in internal processes such as underwriting support, claims documentation and coding, rather than customer-facing use, reflecting a deliberate choice to scale under control rather than scale fast.

Value creation is lagging behind adoption: EIOPA's survey found that most undertakings remain at a proof-of-concept stage. Individual use cases can quickly show productivity gains, but few players have scaled them into enterprise-wide applications. That gap between use-case promise and enterprise-level results is not primarily a technology problem. It is a governance and scaling problem, compounded by uncertainty about how the AI Act and Solvency II will apply to AI embedded in core insurance processes.

Governance is catching up, but from behind: The same survey shows that only 49% of undertakings have a dedicated AI policy today, up from a quarter in 2023. That means roughly half the market is running AI without an AI-specific policy in place.

Regulators have responded, and several obligations already apply: The AI Act imposes binding conformity assessment, monitoring and further obligations on high-risk AI systems, a category that includes AI used for risk assessment and pricing in life and health insurance. Following the AI Omnibus, which entered into force on 27 July 2026, these requirements now apply from 2 December 2027 rather than 2 August 2026. That gives firms more time, though harmonised standards on how the requirements apply in practice are still being developed. EIOPA's own Opinion on AI Governance and Risk Management tells insurers explicitly how their existing Solvency II, IDD, and DORA obligations already apply to AI.

2. Where the pressure is building

Five challenges stand in the way of applying and scaling AI successfully.

Unclear business case and prioritisation: Insurers often cannot agree where AI investment should be concentrated, or whether the benefits outweigh the risks. A lack of prioritisation and unclear value cases are cited more often than budget constraints as the reason AI initiatives stall before they scale. Without an accountable owner to make that case, pilots multiply, but the case for scaling any one of them past proof of concept never gets made. That is the enterprise-level gap described above.

The accountability gap: Most insurers rely on third-party AI models, a sensible commercial choice in most cases, but one where the undertaking stays responsible for a model it may not fully understand. That risk stopped being theoretical in July 2026, when OpenAI disclosed that AI agents running on its models, running with deployment safeguards switched off during an internal cybersecurity test, escaped their sandbox, found previously unknown exploits, and breached the production systems of Hugging Face and other real organisations. Commentators described it as the first true AI safety incident. The episode has sharpened boardroom anxiety about AI systems acting beyond their intended purpose in ways nobody authorised.

A widely reported September 2026 benchmark, RoboHarm, showed the opposite failure. When three AI models were connected to a robot arm and given clearly dangerous instructions, such as placing a compressed air canister on a lit stove, they carried out most of them and refused only a small share.

Hallucinations and inaccurate outputs still top insurers' own list of everyday concerns. Behind all of these sits the same gap. Often nobody inside the organisation both understands the associated risks and is formally accountable for what a model does.

A skills shortage: Regulators themselves flag the scarcity of people who can govern these AI systems as a key constraint, and that shortage will not resolve through recruitment alone in a market where every industry is competing for the same talent.

The trust problem: European consumer research finds that roughly a quarter of consumers remain uncomfortable with insurers' use of AI generally, and views split sharply on AI-driven pricing specifically. Although the AI Act's high-risk category covers life and health insurance, the same questions arise in non-life pricing, particularly motor, where machine learning enables more granular risk segmentation and price optimisation, as EIOPA's 2023 supervisory statement on differential pricing practices recognises. High adoption without visible governance risks turning a regulatory problem into a reputational one, and recent, well-publicised AI failures make that risk harder to dismiss as remote.

Regulatory uncertainty: Large-scale AI deployment is costly, and its return is uncertain for regulatory as well as business reasons. Firms cannot always be confident that an application embedded in underwriting, pricing or claims will meet current and forthcoming requirements, particularly where the AI Act, Solvency II, the IDD, data protection rules and DORA overlap. This is not unique to insurance, but it weighs more heavily in a sector where investments typically take longer to break even. The scaling gap is therefore partly a problem of external, not only internal, governance.

3. What insurers need

Each of these five challenges needs to be addressed.

Business cases and prioritisation need clarity and focus: Estimating the financial and customer impact of AI applications, and prioritising between them, requires analytical rigour and risk-based judgement under uncertainty. That is exactly the toolset actuaries already apply in day-to-day pricing and reserving work.

Accountability needs clear governance and a defined risk appetite: Insurers need clarity on how much risk they are willing to take, where that tolerance ends, and how it varies, since not all AI carries the same risk. A tool used to draft an internal memo is a different proposition from one used to price a policy or settle a claim, which is why a one-size-fits-all policy will not do the job. The relevant methods range from model validation and bias testing to explainability documentation, ongoing monitoring of performance, drift, bias and customer outcomes, and audit trails, the same discipline already applied to actuarial models, extended to a new model class. What is new, as the Hugging Face and RoboHarm cases both showed, is that this discipline now has to cover a model's use and behaviour, not only its internal workings, especially for models an insurer did not build itself.

The skills shortage needs existing actuarial and insurance expertise: Insurers need people who understand the data and statistics behind a model as well as its business and regulatory consequences. This is the combination actuaries have always needed for their core tasks. Actuarial associations are already adding AI, data science, and the relevant regulation to their qualification curricula, and many practising actuaries already have the experience to take on this responsibility.

Trust needs literacy beyond one function: Governance cannot sit only with one function. Underwriters, claims handlers, and customer-facing staff use these tools daily and need enough literacy to know when to escalate, and enough visibility into how decisions are made to keep customer trust intact. Actuaries are well placed to act as that link.

Regulatory uncertainty needs informed dialogue with policymakers: Reducing it requires people who understand both the technical workings of insurance models and the legal requirements that apply to them, and who can translate between the two in dialogue with supervisors and policymakers. Clearer, practical interpretation of how existing and new rules apply to specific insurance use cases would make large-scale AI investment easier to justify.

4. Why actuaries are well placed to help

The profession has adapted to technological change before: When electronic computing arrived in the 1960s, actuaries first used it to automate calculations, then asked how it would change actuarial techniques themselves. In the 1990s, generalised linear models transformed insurance pricing while remaining models that actuaries could understand, challenge and own. Machine learning added predictive power, along with new questions of explainability and validation. Each time, the profession turned technological change into wider professional responsibility.

Foundation models are different in kind, not only in degree: Generative AI and foundation models are general-purpose systems, trained on vast datasets and usually developed and operated by third parties. Their internal behaviour cannot realistically be mastered by an individual user, and is not fully understood even by their developers, as the Hugging Face incident demonstrated in practice. Actuaries may increasingly need to take professional responsibility for the insurance risk arising from models they cannot master internally. The actuarial contribution therefore shifts from model expertise towards model-risk expertise. This means understanding the purpose and materiality of an AI system, challenging its use, assessing the quality and limits of its outputs, setting risk tolerances, defining validation and monitoring, and making sure human judgement and accountability remain in place. This positions actuaries well to connect AI systems to insurance risk, uncertainty, regulation and business consequences, within an established framework of professional standards and accountability, without needing to become AI engineers themselves. It builds naturally on the role Solvency II has given the actuarial function for a decade. Effective AI governance is multidisciplinary: actuaries contribute alongside data scientists, risk and compliance professionals and business leaders.

Reason 1: Actuaries have run regulated model governance for decades: The Hugging Face incident described earlier shows what is at stake when the limits placed on an AI system do not hold: an agent given a narrow task went far beyond it. Actuaries apply the relevant discipline to high-stakes decisions every day: understanding the risk clearly, setting a limit for how much of it is acceptable, and putting controls and sign-off around anything that crosses that limit. Applied to an insurer's own use of AI, that discipline makes risks of this kind far more manageable, and it does not need to be invented. Solvency II has given the actuarial function statutory responsibility for coordinating the calculation of technical provisions, assessing the quality of underlying data, and opining on underwriting policy and reinsurance arrangements since 2016. EIOPA's own reporting goes further than a general statement about the industry. It names actuaries specifically, noting that insurance staff such as actuaries “already count with extensive knowledge and experience on mathematical models” as the natural base from which to extend into new technologies.

There is substantial overlap between EIOPA's AI governance principles and established actuarial practice. Proportionality maps to risk-based technical provisions work, fairness to pricing discrimination testing, transparency to board reporting, human oversight and data governance to the data quality assessment actuaries already perform, and robustness to reserving back-testing. This extends the same discipline to a new, less transparent model class, rather than requiring a new one. The UK's Institute and Faculty of Actuaries reached a similar conclusion independently in its own 2026 AI Manifesto, describing “actuary-in-the-loop” governance as what makes AI systems safer, more reliable, and something a business can stand behind over the long term.

Reason 2: Actuaries bring professional judgement and accountability: The RoboHarm study described earlier shows the other side of the problem: models that could have refused clearly unsafe instructions mostly carried them out. Capability without judgement is not enough. As AI makes analysis cheaper and more widely available, judgement becomes the more valuable resource: framing the problem correctly, choosing which assumptions to trust, understanding uncertainty and materiality, and knowing when an output is not fit for use. This has always been at the heart of actuarial work, including the technical work of building models and preparing data, and it matters more, not less, as AI takes on more of the production. Actuaries work under professional codes of conduct and actuarial standards, such as the European Standards of Actuarial Practice and their national equivalents, backed in most countries by disciplinary schemes, which require them to exercise and document judgement, communicate uncertainty and limitations, and take responsibility for their work. When AI shapes decisions affecting policyholders, insurers need exactly that combination: a named professional who has to judge, document and answer for the outcome, not a system that only executes it.

Reason 3: Actuaries are already shaping the rules, not just complying with them: The Actuarial Association of Europe and its member associations have responded directly to the AI Office's consultation on codes of practice, the Commission's consultation on high-risk AI systems, EIOPA's own consultation on its governance Opinion, and the Digital Omnibus process. At global level, the International Actuarial Association has published important AI-related papers, including its *Artificial Intelligence Governance Framework*. Together, this gives the profession a standing voice in shaping AI governance, in Europe and internationally. For insurers, engaging actuaries and their associations can provide additional insight into regulatory developments and help inform dialogue with policymakers and supervisors. It also gives insurers a channel for reducing regulatory uncertainty, since actuaries and their associations combine technical and regulatory expertise to help policymakers and supervisors clarify how the rules apply in practice.

Reason 4: Actuaries innovate: Actuaries developed credibility theory and stochastic reserving methods, and brought generalised linear models into insurance pricing, decades before “data science” existed as a discipline. That instinct is still visible in how the profession is responding to AI itself. Actuaries' own fear of being displaced by the technology fell from 80% in 2024 to 49% in 2025, according to a survey of US and UK specialty and commercial insurance professionals by the pricing software provider hyperexponential, with a similar drop among underwriters, and insurers are already testing what collaboration looks like in practice. Aviva, for instance, has deployed an AI-assisted “Actuarial Agent” built to support, not replace, its underwriters and actuaries in commercial lines pricing. The claim here is deliberately modest. The capability already sits inside the profession, alongside the governance case, contributing to insurance AI innovation rather than leading it alone.

5. What insurance executives may wish to consider

The case set out above does not call for new structures so much as better use of expertise insurers already have. How far to go will depend on each insurer's size, business mix and AI maturity. The steps below are therefore offered as options to consider rather than a single blueprint. Together, they would place AI governance with professionals who already carry regulatory accountability for models, and help insurers scale AI with confidence. The common thread is that as AI makes models and analysis more abundant, professional judgement, model-risk expertise and accountability become more, not less, valuable.

Where does accountability for AI risk sit?

- Extending the actuarial function's mandate to cover AI in pricing, underwriting and claims, rather than building a parallel function
- Involving actuaries in AI governance roles before making external hires, working alongside the Chief AI or Data Officer rather than in place of them

Where must professional judgement and named accountability remain?

- Defining where professional judgement and named human accountability must remain in AI-supported decisions, with actuaries playing an important role where their expertise is relevant
- Requiring actuarial sign-off wherever AI materially affects technical provisions, pricing or claims assessment processes

How can existing capability be used before creating parallel structures?

- Investing in AI upskilling that builds on the competency frameworks already developed by actuarial associations
- Supporting actuarial-led experimentation with AI where it can add value
- Drawing on actuaries and actuarial associations when engaging with supervisors and policymakers on how AI rules apply in practice

Contributors

Jonas Hirz (main author), Gábor Hanák, Esko Kivisaari, Nicolas Marescaux, Claudio Senatore

AAE Artificial Intelligence and Data Science Working Group

The working group acknowledges and appreciates the review and editorial support of Stephanos Hadjistyllis, AAE Senior Actuary.

This note reflects views from the AAE Artificial Intelligence and Data Science Working Group. Any views expressed are intended to stimulate and inform further discussion and should not be read as a representative or agreed position of the Actuarial Association of Europe.

THE ACTUARIAL ASSOCIATION OF EUROPE

The Actuarial Association of Europe (AAE), founded in 1978 under the name of Groupe Consultatif Actuariel Européen, is the Brussels-based umbrella organisation, which brings together the 38 professional associations of actuaries in 37 countries of the EU, together with the countries of the European Economic Area and Switzerland and some EU candidate countries.

The AAE has established and keeps up-to-date a core syllabus of education requirements, a code of conduct and discipline scheme requirements, for all its full member associations. It is also developing model actuarial standards of practice for its members to use and it oversees a mutual recognition agreement, which facilitates actuaries being able to exercise their profession in any of the countries concerned.

The AAE also serves the public interest by providing advice and opinions, independent of industry interests, to the various institutions of the European Union – the Commission, The Council of Ministers, the European Parliament, ECB, EIOPA and their various committees – on actuarial issues in European legislation and regulation.

AAE | ACTUARIAL ASSOCIATION OF EUROPE

SILVERSQUARE NORTH
BOULEVARD ROI ALBERT II 4
B-1000 BRUSSELS
BELGIUM
+32 2 274 06 61
INFO@ACTUARY.EU
WWW.ACTUARY.EU