



Why involve actuaries in your AI projects?

The AI Manifesto of the Institute and
Faculty of Actuaries (IFoA)

**Involving a Chartered
Actuary makes AI safer,
more robust and more
valuable for the long term**

Contents

Introduction and context	3
What do we mean by “AI”?	3
The skills actuaries already have – and why they matter to your AI project	4
Working with actuaries on AI projects	6
The profession’s commitment	6
Appendix A: What this looks like in practice	7
Appendix B: Frequently asked questions	7

Introduction and context

Artificial Intelligence (AI) already shapes many important decisions in society. The question for organisations and society is no longer whether to use AI, but how to use it well. Doing this requires expertise that brings rigour, clarity and a disciplined understanding of how models behave under uncertainty.

This manifesto is for anyone accountable for, or interested in, the use of AI: executives, product owners, chief risk officers, data and technology leaders, regulators and researchers, whether they are actuaries or non-actuaries. Its aim is to explain how actuarial involvement strengthens your AI project and supports better outcomes for your organisation and the people it serves.

AI needs robust governance, and actuaries have had experience of providing governance under uncertainty since before digital computers existed. AI can feel opaque, and opacity breeds either unwarranted confidence or unwarranted caution. Actuaries help those leading AI projects to understand the real risks clearly and proportionately, so they have the confidence to capitalise on the opportunities.

Actuaries offer a distinctive lens. We are trained to manage uncertainty, balance risk and reward, and apply the professional and ethical standards set out in the Actuaries' Code. The profession is embracing AI with eyes open: positive about its potential, disciplined about its risks, and highly skilled at clarifying complex and potentially unclear trade-offs.

As the IFoA, we believe that actuarial expertise in modelling, risk, ethics and communication is relevant across a wide variety of current AI projects and will become increasingly important as the use of AI systems evolves. For the benefit of society, we encourage those responsible for developing and deploying AI systems to ensure that there is an actuary in the loop.

What do we mean by “AI”?

“AI” has become an umbrella term covering very different technologies, and much confusion in AI projects starts with imprecise language. In this manifesto, and in our work, we distinguish three broad branches:

- **Predictive machine learning (ML)**

Models trained on data to classify, score or forecast, for example pricing, credit, triage and fraud models. Actuaries have been building, validating and governing models of this kind for decades. (We recognise that some commentators exclude this from AI, but actuaries can add considerable value, irrespective of the definition.)

- **Generative AI**

Models, including large language models (LLMs), that produce text, images or code. Their outputs are open-ended, which makes testing and control significantly different from traditional ML.

- **Agentic AI**

Systems that pursue goals with a degree of autonomy, chaining together tools and decisions. These raise distinctive questions of oversight, accountability and failure containment. (We view physical AI as a branch of agentic AI for this purpose.)

The risks, the appropriate controls, and the value an actuary adds differ across these branches, so we encourage every AI project to be explicit about which branch of AI is being used. The IFoA is helping to clarify this terminology, and this manifesto applies to all three branches unless stated otherwise.

The skills actuaries already have – and why they matter to your AI project

Actuaries are professionals trained in modelling, risk and ethics, with a long track record of applying those skills where decisions have serious financial and human consequences. Rather than abstract principles, the case for involving an actuary rests on six core skills the profession has long adopted, and what each one gives you.

What actuaries are trained to do	What that gives your AI project
Build and validate statistical and machine-learning models for high-stakes decisions	Realistic performance claims, robust validation and models that hold up in the real world
Manage large, imperfect datasets under audit and regulatory scrutiny	Trustworthy data foundations, bias identified and managed, faster governance sign-off
Hunt for the ways models fail: edge cases, drift, feedback loops	AI that stays reliable as conditions change, with the right human controls in the right places
Assess risk holistically at enterprise and system level	A coherent view of AI risk across your organisation, adopting systems thinking
Apply professional standards and ethical judgement	Ethics built in from day one, and AI that withstands scrutiny from boards, regulators and the public, reducing both reputational risk and the time to implementation
Translate between technical detail, boardrooms and regulation	Decision-makers who understand what they are approving, and regulation turned into practical steps



Skill 1: We build and validate models for a living – including machine learning

Actuaries are not newcomers to AI. Machine learning (ML) sits on the same statistical foundations as actuarial work, and many actuaries have been building and deploying ML models for a decade in pricing, reserving, capital modelling and beyond. We understand sampling issues, rare events, over-fitting and the gap between performance in testing and performance in the real world. Above all, we are trained to quantify how confident anyone should be in a model's output – and to say so plainly. Actuaries speak up when needed, as early identification of issues is commercially sound as well as professional behaviour.

In practice this means validation that goes beyond headline accuracy: back-testing, challenger models, out-of-sample and out-of-time testing, and sensitivity analysis. For generative AI, where traditional statistical tests do not translate directly, the same mindset applies through curated evaluation sets, structured human review, and monitoring of error and hallucination rates in live use. We challenge unrealistic performance claims before your customers, or your regulator, do it for you.



Skill 2: We manage data quality, bias and governance end-to-end

Every AI system inherits the strengths and flaws of its data. Actuaries are used to working with large, imperfect, transactional datasets subject to audit and regulatory scrutiny. We clarify what data is needed, what quality level is acceptable, and how data should be checked, reconciled and monitored over time. These are core actuarial skills, taught as part of our rigorous professional training.

The same discipline extends to bias. We examine how a model behaves for different groups and sub-groups, test whether training data is representative of the people the system will affect, and help manage the biases we find responsibly through fairness metrics for predictive models, and output audits and red-teaming for generative systems. Good data foundations and honest bias management reduce rework, shorten regulatory review, and produce AI that is fairer and more credible.



Skill 3: We find the ways models fail before the real world does

Actuaries specialise in understanding model limitations. We work with capital models, longevity projections and catastrophe scenarios – all areas where model errors can have far-reaching consequences, both commercially and in their impact on people's lives.

Our role in AI is to explore edge cases, stress scenarios and failure modes. We assess how a system behaves when data drifts, users adapt or feedback loops arise. We help define where weaknesses can be tolerated and where human controls are required. This produces AI systems that remain reliable and well-governed as conditions change.



Skill 4: We see the whole risk picture, across the enterprise and across society

AI risk does not stay neatly inside individual projects. Organisations adopting AI face aggregation questions: many teams relying on the same foundation models or vendors, correlated failure modes, and concentration risk. Actuaries, including those holding the Chartered Enterprise Risk Actuary (CERA) credential, are trained to assess risk holistically at the enterprise level.

For chief risk officers, this means practical help: articulating an AI risk appetite, aggregating AI exposures across the organisation, and reporting to boards in language they can act on. Actuaries are trained to look beyond any single organisation, considering the wider consequences and systemic effects that can arise when many institutions adopt similar systems at once. This perspective aligns with the public interest purpose reflected in the profession's Royal Charter.



Skill 5: We put ethics and the public interest on the table from day one

Actuarial involvement ensures AI is designed with explicit regard to the public interest. Actuaries are bound by professional standards to think beyond immediate commercial outcomes and consider fairness, potential harm and accountability.

In practice, this means building ethics into the project from the start: defining who might be affected, assessing acceptable trade-offs and documenting judgements clearly. Actuaries help teams balance commercial aims with societal expectations, offering constructive solutions rather than imposing limits. This ethical foundation makes the resulting AI more resilient to external scrutiny.



Skill 6: We translate between models and boards, and between regulation and practice

Actuaries are skilled at explaining complex models to non-specialists. As AI attracts increasing attention from boards, regulators and courts, explainability is essential. We help turn technical behaviour into a clear narrative: what the model is designed to do, its strengths and weaknesses, and how it is overseen, so that decision-makers genuinely understand the system being approved.

The same translation skill applies to regulation. Actuaries are not lawyers, but we are demonstrably good at making regulation digestible: working out what accountability, transparency and meaningful human oversight require in practice, and turning them into proportionate steps a project team can deliver.

Working with actuaries on AI projects

Actuaries are used to working in multi-disciplinary teams alongside data scientists, engineers, product owners and legal teams – it is how most actuarial work is delivered. That makes an actuary an easy plugin to an existing AI project: we adapt to your ways of working, bring structure to discussions about risk, performance and fairness, and complement data scientists rather than substitute for them.

We are conscious that professional standards can be caricatured as red tape; our experience is the opposite. Disciplined involvement early in a project is what allows it to safely move fast. Projects with sound validation, honest documentation and clear accountability clear governance more quickly, avoid late-stage rework, and fail in testing rather than in production. Professionalism is what lets an organisation say “yes” to AI with confidence – and protects it from the adverse outcomes that would otherwise force it to say “no”.

Without this type of expertise, AI projects can falter: models that perform well in controlled testing but fail in real-world deployment; biases that go unnoticed until they cause reputational or regulatory harm; AI that is oversold, only to under-deliver when exposed to real users and data; documentation that does not match actual behaviour. Actuarial involvement does not eliminate these risks, but it can materially reduce them.

Nor is this only a message for those buying or building AI. Regulators and researchers need practitioners who combine deep model literacy strong professional and ethical standards and experience of regulated industries. Actuaries are natural collaborators in shaping standards, guidance and research that work in practice.

The profession's commitment

The actuarial profession is developing new skills, research and guidance for AI, building on established strengths while adopting new tools to work confidently with predictive, generative and agentic AI systems. This gives organisations growing access to professionals who understand both AI and long-term risk. The IFoA's commitment to be a hub of innovation and professional excellence are detailed in its 2026-29 Strategy.

AI built with actuarial involvement is more likely to withstand the test of time: better able to cope with changing conditions, easier to audit, and more likely to be accepted by users, employees and regulators.

The IFoA believes that AI can be a powerful force for good when shaped by people trained to balance risk, reward and responsibility. By involving an actuary early in your AI project, you increase the likelihood that the system you build is one you can stand behind over the long term: commercially, professionally and ethically.



The IFoA believes that AI can be a powerful force for good when shaped by people trained to balance risk, reward and responsibility.

Appendix A:

What this looks like in practice

The following short case studies are illustrative, drawn from the kinds of engagement actuaries are already undertaking.

A generative AI customer assistant

An insurer wants an LLM-based assistant to answer customer queries. An actuary designs the evaluation framework: a curated test set of realistic queries, explicit tolerances for error, escalation rules for high-stakes topics, and live monitoring of answer quality. The project clears internal governance at the first attempt because approvers can see exactly what “good enough” means and how it is evidenced.

A machine-learning decision model

A lender deploys an ML model to score applications. An actuary tests behaviour across customer sub-groups, identifies a data artefact that disadvantages one segment, and works with the data science team to correct it. The documentation of that judgement becomes the centrepiece of a constructive conversation with the regulator.

AI risk across an organisation

A Chief Risk Officer faces AI adoption in a dozen business units at once. An actuary helps articulate an enterprise AI risk appetite, maps where the organisation depends on the same foundation models, and builds board reporting that distinguishes experimentation from deployment. The organisation accelerates adoption in low-risk areas precisely because it now knows where the high-risk areas are.

Appendix B:

Frequently asked questions

Why would I pay actuarial fees for this?

Because the expensive AI project is the one that fails late in the process: delayed launches, regulatory findings and reputational damage cost far more than getting the risk and validation work right first time. Actuarial involvement can also be right-sized, from a few days of independent review to an embedded team member.

Aren't actuaries just insurance people?

Insurance and pensions are where the profession grew up, but the skills – modelling under uncertainty, bias management, risk aggregation, communication under scrutiny – apply wherever models drive consequential decisions. Actuaries already work in banking, health, climate, energy and technology.

Won't professional standards slow us down?

The opposite, in our experience. Standards are what make it safe to move quickly: they front-load the questions that would otherwise surface late as governance blockers or public failures. The goal is confident deployment, not caution for its own sake.

Do actuaries really understand generative and agentic AI?

The profession is investing heavily in education, research and practical guidance. More fundamentally, the model validation mindset transfers: knowing what evidence of performance looks like, and what can go wrong between the demo and the real world, matters for every branch of AI.

We already have excellent data scientists – why add an actuary?

Keep the data scientists – actuaries are complements, not substitutes. Data scientists optimise what a system does; actuaries add an independent view of what could go wrong, for whom, and how to evidence that to boards and regulators. The strongest AI teams contain both.



Beijing

Room 512 · 5/F Block A · Landgentbldg Center · No. 20 East Middle 3rd Ring Road
Chaoyang District · Beijing 100022 · People's Republic of China
Tel: +86 10 6611 6828
Email: China@actuaries.org.uk

Edinburgh

Spaces · One Lochrin Square · 92 Fountainbridge · Edinburgh · EH3 9QA
Tel: +44 (0) 207 632 2100

London (registered office)

1-3 Staple Inn Hall · High Holborn · London · WC1V 7QJ
Tel: +44 (0) 207 632 2100

Malaysia

Arcc Spaces · Level 30 · Frankfurt Suite · The Gardens North Tower
Lingkaran Syed Putra · 59200 · Kuala Lumpur
Tel: +60 12 591 3032

Oxford

Belsyre Court · 1st Floor · 57 Woodstock Road · Oxford · OX2 6HJ
Tel: +44 (0) 207 632 2100

Singapore

Spaces · One Raffles Place Mall · #02-01 · Singapore 048616
Tel: +65 9424 9101