LALANI
CONSULTING GROUP

Actuarial Intelligence in the Artificial Intelligence World

Guidance for Actuaries, Employers, and the Profession
2026–2035

A paper for the leadership of the Canadian Institute of Actuaries and the Society of Actuaries

Minaz H. Lalani, FCIA, FSA (1985–2021)

October 2026

LALANI
CONSULTING GROUP

Table of Contents

Executive summary 1

1. Background 4

2. The individual actuary 11

3. The organization 18

4. AI as a risk and a market 29

5. AI by practice area 36

6. The profession: workforce, education, relevance, and standards 44

7. Foresight: scenarios, the task survival model, and indicators to watch 51

8. Conclusion 59

Appendix A. Glossary of AI terms 60

Appendix B. Abbreviations 62

Appendix C. References 64

Disclaimer and Disclosure of AI Use 68

Executive summary

Artificial intelligence is moving from experiment to infrastructure in insurance, pensions, and financial services. This paper, written for the leadership of the Canadian Institute of Actuaries (CIA) and the Society of Actuaries (SOA), sets out what AI means for actuaries and what the profession should do about it. It draws on recent work by actuarial bodies and regulators in Canada, the United States, the United Kingdom, Australia, and Europe, and adds new proposals.

Why now

- OSFI's Guideline E-23 on model risk management, which brings AI and machine-learning models into scope, and the AMF's guideline on the use of AI both take effect on **May 1, 2027**.
- The EU AI Act's obligations for high-risk AI, including life and health insurance pricing, apply from **December 2, 2027**.
- The CIA is reviewing its Standards of Practice for AI and for bias, and its Rules of Professional Conduct. Comments on the bias review are due **October 15, 2026**.

The central argument

AI changes actuarial work at three levels at once — the individual actuary, the employing organization, and the profession — and it is also a new source of risk that actuaries must measure. Existing guidance addresses these levels separately. This paper consolidates them. It concludes that AI is unlikely to make actuaries unnecessary, but it will change what they do: less producing, more supervising, validating, assuring, and advising.

The paper's title is deliberate. **Artificial Intelligence** will become more capable, but it will not carry professional accountability, weigh competing interests or answer to the public. These remain the work of **Actuarial Intelligence**: judgment exercised under uncertainty, in the public interest, by people who stand behind their work.

Key findings

- **Accountability stays with the actuary.** No regulator or professional body accepts "I relied on the AI" as a defence. Actuaries need practical ways to meet that accountability, such as the reviewability test proposed in Section 2.
- **Judgment becomes the core competency.** As AI produces polished answers instantly, the scarce skill is knowing when an answer is wrong (Section 2).

- **Organizations need proportionate governance.** We support the CIA's five-factor AI risk scorecard, propose a severity-dominant rule for combining scores, and propose governing AI that acts on its own as a "digital delegate" (Section 3).
- **Fairness must be tested on outcomes**, not only on the variables used (Section 3).
- **AI is a new insurable peril and a source of systemic risk.** It creates liability exposures, accumulation risk from a few foundation models, and second-order effects on lapse, claims and market assumptions (Section 4).
- **The impact differs sharply by practice area** — fairness in life, health, and P&C; behaviour shifts in group benefits and pensions; correlation in investment, ERM and climate (Section 5).
- **Demand for actuaries holds, but the entry level thins.** US projections show 9% growth in actuarial employment to 2035, while studies show junior employment falling in firms that adopt AI. The risk is fewer experienced Fellows in 10 to 15 years (Section 6).
- **Three moves pay off in every future.** Building AI assurance skills, strengthening reliance and disclosure standards, and protecting the apprenticeship are no-regret actions in all four scenarios we examine. An illustrative task survival model suggests that by 2035, between about half and three-quarters of today's actuarial work would still need a person, with routine work automated first (Section 7).

Ten priority recommendations

Table ES.1. Ten priority recommendations

No.	Recommendation	Lead	Timing
1	Respond to the CIA's consultation on bias in the General Standards	CIA members and practice committees	By October 15, 2026
2	Add AI commentary to the guidance accompanying the Rules of Professional Conduct	CIA	During the 2026 consultation
3	Issue an educational note on reliance on AI, with tiered checks and the reviewability test	CIA, with the SOA	2027
4	Publish a joint AI classification test based on the CIA scorecard, mapped to OSFI E-23, the AMF guideline, and the EU AI Act	CIA and SOA	Before May 1, 2027
5	Prepare for E-23 and the AMF guideline: one AI inventory, digital-delegate controls, validation by risk tier	Employers	By May 1, 2027

No.	Recommendation	Lead	Timing
6	Add at least two AI scenarios to FCT and ORSA	Employers and appointed actuaries	Next FCT and ORSA cycle
7	Co-design a joint AI module and curriculum, and introduce an annual AI and bias CPD hour	CIA and SOA	2027 to 2028
8	Protect the apprenticeship and track the candidate-to-Fellow ratio	Employers; CIA and SOA	Annually from 2027
9	Develop an actuarial AI attestation and recognized AI assurance roles	CIA and SOA	2027 to 2028
10	Commission a member time-use survey, calibrate the task survival model, and maintain an AI foresight dashboard	CIA and SOA	2027

Questions for leadership

- Which future are the CIA and the SOA implicitly planning for, and is it the right one?
- Should standards distinguish AI used as a tool from AI used as a source of judgment?
- If AI reduces entry-level hiring, how will the profession make sure new actuaries still gain the hands-on experience that builds judgment, and who should pay for it: employers, the profession or candidates?
- Should the profession, like the Actuaries Institute in Australia, extend its standards and credentials to non-actuary AI practitioners?

How to read this paper

Each theme sets out what the evidence says, how the issue is likely to develop, our recommended approach for actuaries, employers and the CIA and SOA, and questions for leadership. Tools labelled “proposed by this paper” are our own, built on the sources noted beside them (also included in Appendix C), and are offered for discussion. Key terms are defined in Appendix A and abbreviations in Appendix B.

1. Background

1.1 Why this paper, and why now

Artificial intelligence has moved from experiment to infrastructure in insurance, pensions, and financial services. The question is no longer whether actuaries will use AI, but how the profession keeps its standing as AI does more of the analytical work actuaries have always done.

Three developments make 2026–2027 a turning point:

- **The technology changed in kind, not just degree.** Predictive machine learning has been part of actuarial work for over a decade. Generative AI (since late 2022) and now agentic AI, which can plan and act with limited human involvement, are reaching into drafting, judgment, and decision-making—the core of professional work.
- **Regulation is arriving.** In Canada, OSFI's Guideline E-23 on model risk management takes effect on May 1, 2027, and brings AI and machine-learning models explicitly into scope for federally regulated insurers. In Europe, the AI Act treats AI used in life and health insurance pricing as high-risk; following the Digital Omnibus on AI, those obligations now apply from December 2, 2027. In the United States, the NAIC has issued a model bulletin on insurers' use of AI.
- **The profession is rewriting its own rules.** In 2026, the CIA opened reviews of its Standards of Practice for AI, of bias in its General Standards, and of its Rules of Professional Conduct. The SOA, the IFoA, the Actuaries Institute (Australia), the Actuarial Association of Europe and the International Actuarial Association have each published major AI work in the past 18 months.

Decisions taken in the next two years — on standards, education, and the profession's scope — will shape what an actuary is in 2035. This paper is written to inform those decisions.

1.2 Actuaries and technology: a long relationship

Actuaries have absorbed every major computing advance of the past century. Each one changed the work without shrinking the profession.

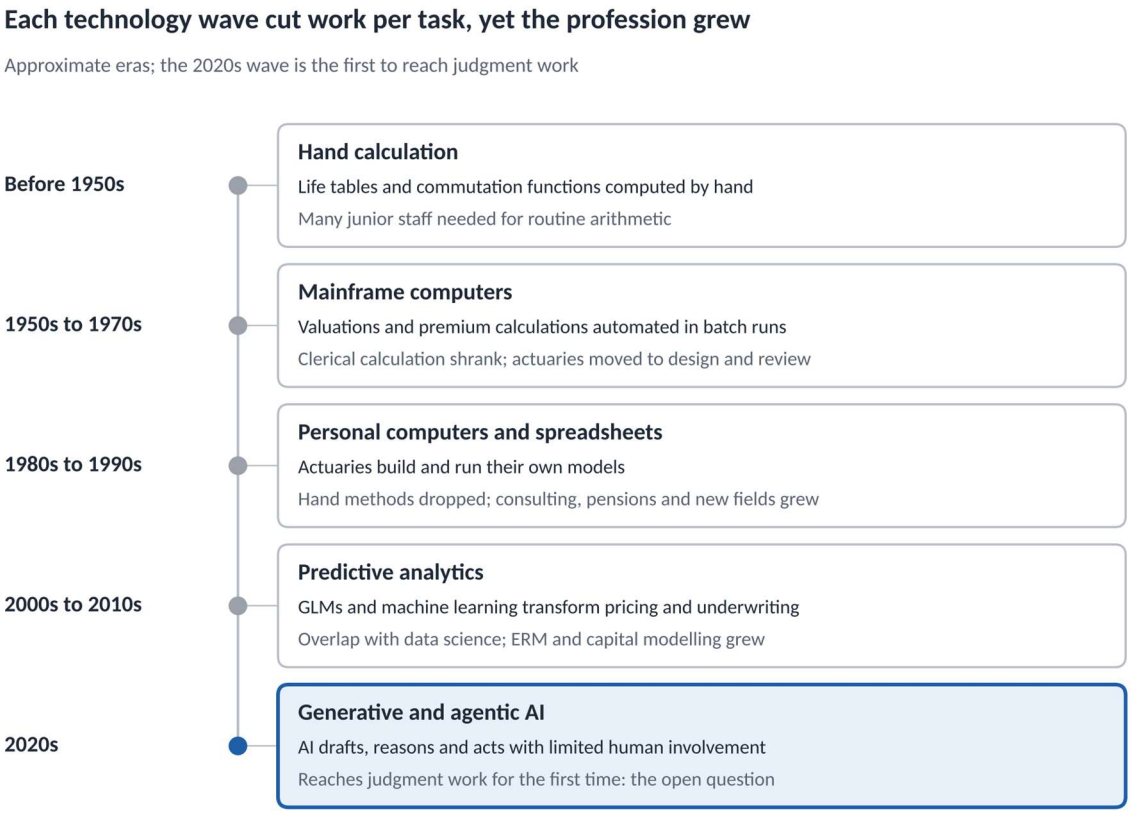


Figure 1. Five waves of technology in actuarial work (approximate eras).

The IFoA's 2025 Presidential Address makes the historical point directly: each step "reduced the number of actuaries needed to carry out any particular task; but as the range of opportunities grew, so did the number of actuaries." Tools that removed calculation freed actuaries for design, judgment, and advice, and opened new fields such as pension consulting, capital modelling, and enterprise risk management.

The current wave may be different in one important respect. Earlier technologies **automated calculation**. Generative and agentic AI reach into **drafting, reasoning, and decision-making** — work that was the actuary's protected core. Whether history repeats depends on whether the profession moves, as it did before, to the new work the technology creates: governing it, validating it, insuring it, and certifying it. This paper addresses that central question.

1.3 What we mean by AI

"AI" covers tools that differ greatly in what they do and the risks they carry. This paper follows the CIA's practice resource *Considerations for AI within an ERM Program*

(September 2026), which treats AI applications as, at their core, a form of model, with added risks from autonomy, opacity, scale, and dependence on third parties.

We use three families throughout:

Family	What it does	Actuarial examples	Main new risk
Predictive AI (machine learning)	Learns patterns from data to estimate an outcome	GLMs, gradient boosting and neural networks for pricing, underwriting, fraud, and reserving	Opacity, bias through proxy variables, drift
Generative AI	Produces text, code, summaries, or images from instructions	Drafting reports, summarizing medical files, writing model code, research support	Plausible but wrong output; hidden reliance
Agentic AI	Plans and carries out multi-step tasks, using tools and acting with limited human involvement	Running a valuation end to end, triaging, and settling simple claims, monitoring models	Actions taken without review; unclear accountability

The boundary between a traditional actuarial model and an "AI system" is not sharp. A GLM is a statistical model; a GLM that an agent fits, tunes, and deploys automatically is part of an AI system. For that reason, this paper classifies AI by **what it does and how much it decides**, rather than by the technique inside it.

The table above describes the three families. Other key terms used throughout the paper are defined in the box below. Where regulators define a term differently, a footnote explains the difference, and Appendix A gives a fuller glossary.

Key terms used in this paper

AI system — A machine-based system that, with some degree of autonomy, infers from its inputs how to generate outputs such as predictions, content, recommendations, or decisions.[1]

Large language model (LLM) — A generative AI model trained on very large amounts of text, the basis of chatbots and AI writing tools.

Foundation model — A large, general-purpose model adapted to many uses. A handful underpin most generative AI, which concentrates risk on a few providers.

Hallucination — Fluent, confident output that is false or invented, such as a fabricated citation or figure.

Explainability — The ability to explain why a model produced a given output, in terms a user or regulator understands.

Model drift — A decline in a model's performance as data or conditions move away from those it was trained on.

Proxy discrimination — Unfair outcomes that arise through variables standing in for a protected characteristic, such as postal code.

Human-in-the-loop — A person reviews, approves, or can override AI output before it takes effect.

1.4 Where AI stands in insurance and pensions today

AI is now embedded across the insurance value chain, but its value is uneven, and its governance lags its use.

Breadth of use. The CIA's ERM resource describes AI across pricing, underwriting, claims, marketing, finance, and fraud detection. In P&C, it supports claims triage, fraud detection, and document summarization. Across practice areas, it supports asset-liability management, hedging, and scenario testing. Pension funds use it for investment analysis and fraud analytics. In the United States, the NAIC's Big Data and AI Working Group has surveyed insurers by line. In its surveys, about 88% of private passenger auto insurers and about 70% of home insurers responding said they use, plan to use or plan to explore AI or machine learning.[2]

---

[1] The EU AI Act (Regulation (EU) 2024/1689, Article 3(1)) defines an "AI system" as "a machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments." OSFI's Guideline E-23, Model Risk Management (2027), takes a broader approach: it treats as a "model" any method that processes input data to generate results, whether or not AI is involved. A traditional GLM is therefore a model under E-23 but may fall outside the EU definition of an AI system. Many other laws avoid a precise definition of AI and treat AI tools as forms of modelling (CIA, Considerations for AI within an ERM Program, 2026, Section 2).

[2] National Association of Insurance Commissioners, Big Data and Artificial Intelligence (H) Working Group, summary of private passenger auto and home insurance surveys (August 2023).

Uneven value. The SOA's expert panel on life underwriting (July 2026) found AI is already producing value, but unevenly, depending on carrier maturity, data readiness, workflow design, and staff skill. Its advice was to start from the business problem, not the technology.

Rising concern. In the IFoA and London Foundation for Banking and Finance survey (May 2026), 70% of senior financial services practitioners agreed AI risks are among the greatest facing their sector over the next five years, and 75% said those risks had grown substantially since generative AI arrived. The top three risks were cyber threats, misleading outputs, and knowledge gaps.

Pressure on skills. KPMG's 2025 Insurance CEO Outlook (January 2026) found that 72% of insurance CEOs see AI workforce readiness and upskilling as a top constraint on growth, and 79% say AI is changing the skills required for entry-level roles[3]. In Australia, about one-third of Actuaries Institute members now specialize in data science.

The picture is one of fast, broad adoption with governance, skills and professional standards still catching up. That gap is where the actuarial profession can add most value — or lose ground to others.

1.5 How the profession has responded so far

Every major actuarial body has engaged with AI, but each has emphasized a different level: the individual, the organization, or the profession.

Jurisdiction	Key work	Main emphasis
Canada (CIA, OSFI)	Considerations for AI within an ERM Program (Sept 2026); Notice of Intent to review the Standards of Practice for AI (May 2026); bias review of the General Standards (comments due Oct 15, 2026); OSFI Guideline E-23 (effective May 1, 2027)	The organization: governance, validation, model risk
United States (SOA, CAS, Academy, ASB)	Actuarial Mindsets for Leading in the AI Era (Jan 2026); SOA AI research program and strategic priority to "Leverage AI"; AI the largest shared priority of the Academy's practice councils in 2026; ASOP revisions under way	The individual: mindset, judgment, accountability

[3] KPMG, 2025 Insurance CEO Outlook (January 2026).

United Kingdom (IFoA)	It's Still Not Magic with LFBF (May 2026), a nine-risk framework; AI Risk Alert (2023)	Systemic and sector risk
Australia (Actuaries Institute)	Strategic Plan 2025–2027: Reach and Impact in the Age of AI; compulsory Data Science Principles subject	The profession: education, brand, wider data professions
Europe (AAE)	Discussion paper on the EU AI Act (Mar 2025); support for clarifying that GLMs are not automatically high-risk AI (Apr 2026)	Regulation and classification
International (IAA)	AI Task Force (2024); papers on governance, testing and documentation (Mar 2026)	Global professional framework

Common ground is strong. All bodies agree the actuary remains accountable regardless of the tools used, that AI should be governed as a model with added risks, and that education must change. Where they differ is ambition. Australia is openly reshaping the profession's scope; most others are adapting existing frameworks.

1.6 The gap this paper addresses

The existing work is sound but fragmented. Four gaps stand out.

1. **No single view across the three levels.** Guidance for the individual actuary, the employing organization and the profession has been developed separately. A change at one level — for example, firms hiring fewer juniors — has effects at the others, such as fewer experienced Fellows a decade later. These links are not yet addressed.
2. **Little quantification.** Current guidance describes AI risks but rarely measures them. There is no agreed way to put AI scenarios into capital tests, to price AI as an insured peril, or to estimate how much actuarial work will change.
3. **Little forward-looking planning.** Most guidance responds to today's tools. Agentic AI, AI-reviewed regulatory filings, and AI-driven customer behaviour are discussed rarely, if at all.
4. **Guidance by practice area is thin.** The syllabus is organized by Life, Group Benefits, Pensions, P&C, Investment and ERM, but most AI guidance is general. Practitioners need to know what changes in their own area.

This paper aims to close these gaps with clear recommendations, while being open about where evidence is still limited.

1.7 Scope, audience and approach

Audience. The paper is written for the boards, education committees and standards bodies of the CIA and the SOA. It should also be useful to employers of actuaries and to regulators.

Scope. It covers actuarial work in Canada and the United States, drawing on evidence from the United Kingdom, Australia, and Europe. It addresses all main practice areas: life and annuities, group benefits and health, pensions and retirement, P&C, investment, ERM and climate.

Approach. The paper combines three methods:

- **Evidence review** of published work by actuarial bodies, regulators, and industry in five jurisdictions.
- **Anticipatory thinking**: four scenarios to 2035, built on two key uncertainties — the pace of AI capability and the strength of governance — and a set of early-warning indicators.
- **Actuarial methods applied to the profession itself**: a task-level survival model that treats each actuarial task's automation as a hazard, to estimate where work will shrink and grow.

Structure. After this background, the paper is organized in six parts:

1. The individual actuary
2. The organization
3. AI as a risk and a market
4. AI by practice area
5. The profession: workforce, education, relevance, and standards
6. Foresight: scenarios, the task survival model, and indicators to watch

Each theme sets out the evidence, how it is likely to develop, recommended guidance for actuaries, employers and the CIA/SOA, and questions for leadership.

Status and limitations. This is a discussion paper. It is not an actuarial opinion and does not form part of any standard of practice. Recommendations have been considered against

the CIA Standards of Practice and Rules of Professional Conduct, both under review in 2026. The field is moving quickly: the evidence is current to October 2026; Figures are cited to original sources where available; two items rely on secondary reports and are noted as such in the references.

Proposed tools. Tests, tables, and frameworks labelled "proposed by this paper" are this paper's own. Each builds on the sources noted beside it and is offered for discussion, not as settled practice or a standard of practice.

Abbreviations. Abbreviations used in this paper are spelled out in Appendix B, together with the regulatory guidelines and standards referred to.

2. The individual actuary

AI changes what an actuary does. It does not change who answers for the work. This section sets out what every practising actuary needs: a clear rule for accountability when AI contributes to the work, the habits of mind that make that accountability real, and a sensible standard for telling users when AI shaped the result.

Each subsection follows the same pattern used throughout the paper: what the evidence says, how the issue is likely to develop, our recommended approach, and the questions only CIA and SOA leadership can settle.

2.1 The signatory principle: accountability for AI-assisted work

The principle. An actuary who signs a piece of work is responsible for all of it, whatever tools produced it. This is not new. It already applies to spreadsheets, vendor models and work prepared by junior staff. AI tests the principle more strongly because its output is fluent, fast, and often hard to trace to its sources.

What the profession says now

- **Society of Actuaries.** The SOA's expert panel on AI readiness (January 2026) concluded that responsible work must let another actuary judge whether it is sound, and "cannot rely on an actuary saying, 'I relied on the AI.'" Responsibility lies with the actuary signing the work, not with the model.
- **Canadian Institute of Actuaries.** The Actuarial Standards Board's Notice of Intent (May 26, 2026) is reviewing whether current standards adequately address an actuary's responsibilities when relying on AI tools and third-party outputs. It also covers data quality, representativeness and bias, model validation, explainability, documentation, and alignment with OSFI's risk-based approach to model oversight in Guideline E-23.

- **Actuarial Association of Europe.** In its response on the EU AI Act, the AAE stated that actuaries, insurers, and pension providers bear full professional responsibility for their decisions and outputs, regardless of the AI tools they employ.
- **United States.** ASOP No. 56, *Modelling*, already sets expectations when an actuary relies on a model developed by others. The Actuarial Standards Board is also revising ASOP No. 41, *Actuarial Communications*; its third exposure draft (March 2026) expands disclosure guidance.

How existing Canadian requirements already apply

No new rule is needed to establish accountability; the existing framework already covers AI. What is missing is guidance on how to meet it in practice.[4]

Table 2.1. How existing CIA requirements apply to AI-assisted work

Requirement	Source	What it means when AI contributes
Qualification	Rules of Professional Conduct	Use AI only on work you are qualified to do yourself and able to check. AI does not extend an actuary's competence.
Appropriate methods and assumptions	Standards of Practice, General Standards (Part 1000)	A method or assumption suggested by AI must be justified on its own merits, as if the actuary had chosen it.
Use of models	General Standards	AI output is model output. The actuary must understand the tool well enough to know its limits, validate the result and document reliance.
Reliance on others' work and data	General Standards	A third-party AI tool is treated like any external source: disclose the reliance and the extent of review.
Reporting	General Standards	The report states purpose, methods, assumptions, and limitations, including any material role played by AI.
Confidentiality and control of work product	Rules of Professional Conduct	Client or employer data must not be entered into AI tools that retain or train on it without authority. AI-drafted text must be reviewed before release.

Looking ahead

By 2035, AI agents are likely to prepare complete first drafts of valuations, rate filings, and actuarial reports. The actuary's role will shift from producing the work to supervising and

[4] Specific section and rule numbers are not cited here because the CIA Standards of Practice and Rules of Professional Conduct are both under review in 2026. Readers should refer to the current consolidated texts on the CIA website.

signing it. Regulators are also beginning to use AI to review filings, so actuarial work will increasingly be read by machines as well as people.

The main risk in this shift is **automation complacency**, a well-documented effect in human-factors research: people who supervise reliable automation tend, over time, to check it less carefully.[5] As volume rises and review time per item falls, sign-off can become nominal without anyone deciding it should. Aviation and medicine manage this risk by design, not by good intentions, and actuarial practice should do the same.

Recommended approach

We propose two practical tools. The first is a short test every actuary can apply before signing. The second scales the checks to the risk of the AI use, so that low-risk uses are not burdened and high-risk uses get real scrutiny.

The reviewability test — four questions before signing (proposed by this paper)

1. **Reproducible.** Could a qualified peer reproduce the key results without the AI tool, or from a documented version of it?
2. **Explainable.** Can I explain each material method and assumption in my own words, and justify it independently of the AI?
3. **Independently checked.** Have I tested the output against something independent — a simpler model, prior results, or a reasonableness range?
4. **Recorded.** Is the AI's role on file — which tool and version, when, and what it did?

If any answer is no, the work is not ready to sign.

Builds on: *the SOA expert panel on AI readiness (2026) on reviewability and explainability; OSFI Guideline E-23 and ASOP No. 56 on independent validation; OSFI E-23 and the CIA ERM resource (2026) on model inventories and rationale logging.*

Table 2.2. Proportionate checks by level of AI involvement (proposed by this paper)

Tier	Example AI use	Minimum checks before reliance
Low	Editing, formatting, translating the actuary's own text; summarizing public documents	Read the output. No record needed beyond the normal working file.
Moderate	Drafting report text; writing or debugging model code; research support	Line-by-line review. Check every figure, citation, and quotation against its source. Record the tool and its use.
High	AI suggests or produces methods, assumptions, or results relied on in a signed opinion	Independent recalculation or challenger model; documented justification; peer review; disclosure in the report.

[5] Parasuraman, R. and Manzey, D. H. (2010). "Complacency and Bias in Human Use of Automation: An Attentional Integration." Human Factors, 52(3), 381–410.

Tier	Example AI use	Minimum checks before reliance
Critical	AI acts without human approval — sets a price, approves a claim, books a reserve	Governed as an organizational system (Section 3), with a written mandate, limits, and monitoring. The actuary opines on the controls and aggregate results, not on each action.

Builds on: *the risk-based approach of OSFI Guideline E-23 and the five-factor risk scorecard in the CIA ERM resource (2026).*

For actuaries: apply the reviewability test to all signed work, and keep a simple AI log for each engagement.

For employers: keep a register of approved AI tools and their permitted uses; budget review time explicitly, so that volume targets do not quietly erode review.

For the CIA and SOA: issue an educational note on reliance on AI that adopts proportionate, tiered checks; and coordinate wording with the US Actuarial Standards Board so that actuaries working across borders face one set of expectations.

Questions for leadership

- Should the standards distinguish AI used as a **tool** (like a spreadsheet) from AI used as a **source of judgment** (like an external expert), with different reliance rules for each?
- Should actuaries signing high-risk work be expected to confirm, in the report, that the work meets a reviewability standard?

2.2 The AI-ready actuary: judgment as the core competency

The issue. As AI produces polished answers instantly, the scarce skill is no longer producing an answer. It is knowing when an answer is wrong. That depends as much on mindset and habit as on technical knowledge.

What the evidence says

- The SOA panel defined AI readiness by mindset, not coding skill. AI-ready actuaries treat AI as "informative rather than authoritative," apply "disciplined skepticism without becoming obstructionist," show intellectual humility, and interpret AI output as a probabilistic signal rather than an answer.
- The panel described large language models as "a very good productivity tool, but a very bad teaching tool," and argued that new actuaries need "time and breathing room" for first-principles work — "the Excel actuary kind of needs to come back."

- It also observed a learning curve: early users tend to overestimate what AI can do, and genuinely AI-ready people move to a calibrated view of its strengths and limits.
- The CIA's ERM resource independently lists the **erosion of critical thinking** among staff as a human-resources risk of AI. Two bodies, working separately, have flagged the same concern.
- The Actuaries Institute's 2025–2027 strategy stresses executive and "soft" skills — leading multidisciplinary teams and influencing decisions — alongside technical depth.

Looking ahead

AI is weakest where data is thin, conditions are changing, or the risk is new: emerging perils, regime shifts, long-tailed liabilities, and judgment-based distributions. These are precisely the actuary's home ground. The actuaries who thrive will be those who can supervise AI on routine work and exercise judgment where AI cannot — and who can explain the difference to a board.

Recommended approach

We propose that the profession define AI readiness as a four-level competency ladder. It gives employers a shared language for roles, gives the education system clear targets, and lets CPD tie to a level rather than hours alone.

Table 2.3. AI-readiness competency ladder (proposed by this paper)

Level	What the actuary can do	Typical stage	How it is developed
1. Aware	Understands the main AI families, what each does and how each fails; uses approved tools for low-risk tasks	Student	Foundation module, assessed without AI
2. Practitioner	Uses AI productively and checks its output reliably; applies the reviewability test	Associate	Applied modules, assessed with AI permitted
3. Supervisor	Designs workflows in which AI does the work and people check it; reviews others' AI-assisted work	Fellow, manager	CPD and supervised practice
4. Governor	Sets AI policy, validates high-risk AI, advises boards and regulators on AI risk	Appointed actuary, chief actuary, CRO	Specialist credential (Section 6)

Builds on: *the SOA expert panel's description of how AI readiness develops (2026) and the Actuaries Institute's emphasis on leadership and multidisciplinary skills (Strategic Plan 2025–2027).*

Four habits support every level:

- **Break the model.** Test boundary cases and ask what the AI could not have seen.

- **Two sources for every number.** Confirm any figure that matters against a second, independent source.
- **Signal, not answer.** Treat AI output as one input to a judgment, never as the judgment.
- **First principles first.** New actuaries do core calculations by hand before automating them.

Questions for leadership

- How should qualification assess **judgment** — including the ability to critique AI output — rather than knowledge alone?
- Should AI-readiness levels be recognized formally, for example in CPD requirements or role descriptions for appointed actuaries?

2.3 Disclosure: telling users when AI shaped the work

The issue. Users of actuarial work — boards, regulators, courts, and plan members — increasingly want to know whether AI played a part. Too little disclosure undermines trust; too much turns into boilerplate that nobody reads.

What the evidence says

- Neither the CIA Standards of Practice nor the US ASOPs yet address AI-assisted actuarial communications specifically. The revision of ASOP No. 41, now in its third exposure draft (March 2026), is the closest development.
- Wider regulation is moving toward transparency. The EU AI Act's transparency obligations for certain AI interactions and AI-generated content were not deferred by the 2026 Digital Omnibus, unlike its high-risk obligations.
- Since January 2026, US insurers have been able to add optional generative-AI exclusions to commercial general liability policies (Section 4.1). If this extends to professional liability, an actuary relying heavily on AI may carry uninsured exposure.

Looking ahead

Disclosure of AI use is likely to become routine in actuarial reports, much as disclosure of model limitations and climate assumptions has. The real debate will be about materiality: routine drafting help should not need disclosure, but AI that shaped a method, assumption or result should.

Recommended approach

We recommend a materiality-based approach that mirrors how actuaries already treat reliance on others.

Table 2.4. Materiality-based disclosure of AI use (proposed by this paper)

Type of AI use	Disclose?	Suggested approach
Editing, formatting, or translating the actuary's own text	No	None needed.
Research or summarizing source material, checked by the actuary	Usually no	Keep a record on file.
Drafting substantial parts of the report	Yes, briefly	For example: "Parts of this report were drafted with the assistance of a generative AI tool and reviewed by the actuary."
AI suggested or produced a method, assumption, or result	Yes, specifically	Name the tool, describe its role, and state how the result was validated.
An AI system makes or drives decisions the report relies on	Yes, in full	Describe the system, its governance, its validation, and its known limitations.

Builds on: *existing practice for disclosing reliance on the work of others under the CIA General Standards.*

Two related precautions apply now, before any new standard: check professional liability cover for AI exclusions and agree with clients or employers in writing which AI tools may be used on their work and how their data will be protected.

Questions for leadership

- Where should the materiality threshold sit, and should the CIA and SOA publish standard disclosure wording?
- Should the profession engage with professional liability insurers now, so that AI-assisted actuarial work remains insurable?

2.4 Section 2 at a glance

Table 2.5. Summary of recommendations for the individual actuary

Audience	Recommendations
Actuaries	Apply the reviewability test before signing. Scale checks to the tier of AI involvement. Disclose material AI use. Build the four habits. Check professional liability cover.
Employers	Keep a register of approved AI tools. Budget review time explicitly. Protect first-principles training for new actuaries. Agree AI use and data protection with clients in writing.

Audience	Recommendations
CIA and SOA	Issue an educational note on AI reliance with tiered checks. Adopt the AI-readiness competency ladder. Publish standard disclosure wording. Coordinate with the US Actuarial Standards Board and engage professional liability insurers.

3. The organization

Most actuaries use AI inside an organization: an insurer, a pension plan, a consultancy, or a regulator. That organization decides which tools to use, how to approve and check them, and how to measure their risks and hold capital against them. This section turns to those decisions.

The starting point is the CIA's practice resource *Considerations for AI within an ERM Program* (September 2026). It treats AI applications as "at their core, a form of model," and AI risk as an extension of model and operational risk, "with additional considerations for autonomy, opacity, scale and third-party dependencies." It also warns that "AI governance is more than a model validation challenge; it is an enterprise risk, technology, compliance and strategy challenge." We agree, and build on that framework in five areas.

Two Canadian requirements take effect on the same day, **May 1, 2027**: OSFI's Guideline E-23 on model risk management for federally regulated insurers, and the Autorité des marchés financiers' *Guideline for the Use of Artificial Intelligence* for insurers and other financial institutions authorized in Quebec. Much of this section is about preparing for that date.

3.1 Classifying AI use: where a model ends, and AI begins

The issue. Organizations cannot govern AI in proportion to its risk until they can sort one use from another. Regulators draw the boundary in different places, and none of them draws it by the technique used.

What the evidence says

- **European Union.** The AI Act lists AI systems used for risk assessment and pricing of natural persons in life and health insurance as high-risk (Annex III). Following the 2026 Digital Omnibus, those obligations apply from December 2, 2027.
- **Actuarial Association of Europe.** In April 2026, the AAE supported a call for the European Commission to clarify that generalized linear models and similar actuarial models, already governed by mature actuarial frameworks, should not automatically be classed as high-risk AI.

- **OSFI.** Guideline E-23 defines a model broadly as any method that processes input data to generate results. Traditional actuarial models and AI models are both in scope; what changes is the intensity of governance, which follows each model's risk rating.
- **AMF.** The Quebec guideline requires each institution to keep a centralized directory of its AI systems, assign each a risk rating, and review the ratings periodically.
- **CIA.** The ERM resource scores each AI use case from 1 to 5 on five factors — impact severity, autonomy and human oversight, complexity and explainability, data sensitivity, and deployment scale — and combines them into four tiers: low, moderate, high, and critical. It does not say how the scores are combined.

Table 3.1. The CIA's five-factor AI risk scorecard

Factor	What is assessed	Score 1 (lowest risk)	Score 5 (highest risk)
1. Impact severity	Financial loss, customer harm, and regulatory and reputational risk if the AI fails or has downstream effects	Negligible	Systemic or critical impact
2. Autonomy and human oversight	Degree of automation, level of human review and whether decisions can be reversed	Advisory, with mandatory human approval	Fully autonomous, with limited oversight
3. Complexity and explainability	Model complexity (for example, deep neural networks compared with GLMs), interpretability and stability	Simple and transparent	Highly complex and opaque
4. Data sensitivity	Personal information, protected health information and intellectual property	Public or low sensitivity	Highly sensitive or regulated
5. Deployment context and scale	Pilot or organization-wide use, internal or external exposure, and how critical the process is	Limited pilot	Enterprise-wide and customer-facing

Source: *CIA, Considerations for AI within an ERM Program (September 2026), Section 3, "Risk assessment." Each factor is scored from 1 to 5, and the scores combine into four tiers: low, moderate, high, and critical.*

Looking ahead

Classification by technique will not last. A GLM fitted, tuned, and deployed automatically by an AI agent is technically still a GLM, but it carries the agent's risk. Within a few years, most models of any kind will be built or maintained with some AI assistance. Classification by **what the system does and how much it decides** will remain meaningful.

Recommended approach

We support the CIA's five-factor scorecard as the common Canadian basis for classification. We propose one addition: a clear rule for combining the scores. A simple average can let a severe, autonomous use hide behind low scores on other factors. Insurance practice already knows the remedy — severity dominates.

Table 3.2. A severity-dominant rule for combining the CIA scorecard (proposed by this paper)

Step	Rule
1. Base tier	Average the five factor scores. Below 2 = Low; 2 to under 3 = Moderate; 3 to under 4 = High; 4 or more = Critical. These thresholds are illustrative and should be calibrated by each organization.
2. Severity override	Any single factor scored 5 makes the use case at least High.
3. Autonomy override	Impact severity of 4 or more combined with autonomy of 4 or more makes the use case Critical.
4. Review	Re-score whenever the use case, the model, or the vendor changes, and at least annually.

Builds on: *the five-factor risk scorecard and four risk tiers in the CIA ERM resource (2026); severity-led thinking common in insurance risk assessment.*

Two examples show the effect:

Use case	Scores (impact, autonomy, complexity, data, scale)	Average	Final tier
Generative AI tool summarizing public regulatory documents for internal use	2, 2, 4, 1, 3	2.4	Moderate (no override)
Pilot model that automatically declines small claims without review	4, 5, 3, 2, 1	3.0	Critical (autonomy override); the average alone gives High

For organizations: keep one inventory of AI uses that serves E-23, the AMF guideline and internal ERM at once, scored on the same five factors.

For the CIA and SOA: publish a joint, jurisdiction-neutral classification test, mapped to E-23, the AMF guideline, the EU AI Act and US state requirements, so that cross-border organizations score once.

Questions for leadership

- Should the Canadian profession seek, as the AAE has in Europe, explicit recognition that well-governed traditional actuarial models are not "high-risk AI"?
- Should the CIA recommend standard thresholds for the scorecard, or leave calibration to each organization?

3.2 Governance and validation: from model to digital delegate

The issue. Today's governance frameworks were built for models that calculate and people who decide. AI agents increasingly do both. The CIA's ERM resource asks the right question — when an AI application should be managed within model risk frameworks and when "more akin to a staff member or external advisor" — but does not answer it.

What the evidence says

- **Three lines of defence.** The CIA resource assigns the first line (business and technology) ownership of AI outcomes; the second line (risk and compliance) challenges, risk ratings, the AI inventory, model validation, and bias testing; and the third line (internal audit) audits across the AI lifecycle.
- **Lifecycle.** OSFI E-23 sets expectations at each stage of a model's life, from design and data through review, approval, deployment, monitoring, and decommissioning. The AMF guideline sets out a comparable lifecycle and makes the board responsible for a culture of responsible AI use and for the board's collective competence on AI risk.
- **Validation.** The CIA resource sets out eight levels of validation: basic functional, performance, robustness, bias and fairness, explainability and interpretability, regulatory and compliance, operational, and post-deployment. It suggests organizations deepen validation as their AI adoption matures.
- **Monitoring.** It describes three pillars: ongoing testing with thresholds and tier-based frequency; vendor and third-party oversight; and drift detection with trigger-based retraining.

Looking ahead

AI that acts — sets a price, approves a claim, rebalances a hedge or sends a customer letter — will not fit frameworks designed for tools that only produce numbers for a person to use. The controls organizations already use for people who act on their behalf—mandates, signing authorities, supervision, and performance review—are a better fit.

Recommended approach

We propose answering the CIA's question directly, with two governance tracks. **Advisory AI**, whose output a person reviews before acting, stays under model risk management. **Acting AI**, which takes actions itself, is governed as a **digital delegate**: model risk management plus the controls used for delegated human authority.

Table 3.3. Two governance tracks for AI (proposed by this paper)

Element	Advisory AI: model risk track	Acting AI: digital-delegate track
Framework	OSFI E-23 model lifecycle	E-23 lifecycle plus delegated-authority controls
Approval	Model approval before use	A written mandate: what the AI may do, for which business, within which limits
Limits	Restrictions on permitted uses	Authority limits like signing authorities — for example, a maximum claim amount or price change
Oversight	Periodic validation	A named human supervisor, regular sampling of actions, and defined escalation triggers
Performance	Accuracy, calibration, and drift metrics	A periodic performance review: error rates, overrides, complaints, and outcomes
Stopping	Decommissioning	A tested off switch and a manual fallback process
Records	Model documentation	A tamper-proof log of every action, with its inputs and rationale

Builds on: *the CIA ERM resource's question of when AI should be managed as a model or as a staff member (2026); OSFI Guideline E-23 lifecycle expectations; existing delegated-authority controls such as signing authorities.*

We also propose one change to how validation is scaled. The CIA resource links validation depth to an organization's overall AI maturity. We recommend linking it to the **risk tier of each use case** instead. A small insurer with a single AI pricing model still needs fairness and explainability testing for that model.

Table 3.4. Validation depth by risk tier, using the CIA's eight validation levels (proposed by this paper)

Validation level	Low	Moderate	High	Critical
Basic functional	Required	Required	Required	Required
Performance	Required	Required	Required	Required
Regulatory and compliance	Required	Required	Required	Required
Bias and fairness	If people are affected	If people are affected	Required	Required
Explainability	Not required	Recommended	Required	Required

Validation level	Low	Moderate	High	Critical
Robustness	Not required	Recommended	Required	Required
Operational	Not required	Required	Required	Required
Post-deployment monitoring	Annual review	Periodic	Continuous	Continuous, with alerts

Builds on: *the eight validation levels and tier-based monitoring frequency in the CIA ERM resource (2026); the risk-based approach of OSFI Guideline E-23.*

For organizations: actuaries should design or approve the validation of any AI that affects pricing, reserving, capital, claims, or member benefits. That is where actuarial expertise in model risk is strongest and where the consequences of error are largest.

For the CIA and SOA: develop guidance on validating AI models in actuarial use, including agentic systems, and consider whether AI validation should become a recognized area of actuarial expertise.

Questions for leadership

- Should "AI validation actuary" become a recognized specialty, with its own credential?
- Who should own an AI agent's mandate — the business, the chief actuary, or the chief risk officer?

3.3 Fairness, bias and explainability

The issue. Insurance is built on differentiating by risk. AI makes it possible to differentiate on far more variables, some of which may stand in for characteristics that must not be used. The task is to keep legitimate risk-based pricing while preventing unfair discrimination and explaining the difference.

What the evidence says

- **CIA standards.** The CIA is consulting on bias considerations in its General Standards. Comments are due October 15, 2026; an exposure draft is targeted for January 31, 2027, and final standards will be effective no later than May 1, 2027. The CIA is coordinating the work with its AI working group.
- **CIA ERM resource.** Bias can persist even when personal information is removed, when models are trained on historical human decisions, or when proxy variables such as postal code are used to infer protected attributes. Its fairness validation refers to measures such as disparate impact and equalized odds.

- **AMF.** The Quebec guideline expects institutions to tell clients when they are interacting with an AI system, offer access to a person, and give a simple, clear explanation to clients subject to decisions made by or with the help of AI.
- **Quebec privacy law.** Quebec's private-sector privacy legislation already requires organizations to inform a person when a decision about them is based exclusively on automated processing, and to explain the main factors on request.[6]
- **Federal and provincial.** OSFI and the Financial Consumer Agency of Canada emphasize explainability, data quality, governance, and ethics. Provincial rate regulators review rating variables, and human rights legislation applies to AI-driven decisions as to any other.

Looking ahead

Fairness will move from a pricing question to a whole-lifecycle question covering marketing, underwriting, claims, and customer communication. Regulators are likely to ask for evidence from **outcome testing** — how different groups actually fare — rather than a list of variables that were excluded.

This creates a practical problem, particularly acute in Canada. Most insurers do not collect characteristics such as ethnicity, so they cannot directly test outcomes for those groups. Options include estimating group membership statistically, collecting data voluntarily with consent, or testing through an independent third party. Each raises privacy and accuracy questions that the profession should help resolve.

Recommended approach

We propose a three-step approach for any AI model that affects people (proposed by this paper):

- **Justify the inputs.** Document why each rating or decision variable is a legitimate indicator of risk and look for variables that act mainly as proxies.
- **Test the outcomes.** Compare prices, decisions, and errors across groups, after allowing for genuine risk differences, using one or more fairness measures chosen in advance.
- **Explain and remediate.** Explain any individual decision in plain language, and act when testing reveals an unjustified difference.

[6] Act respecting the protection of personal information in the private sector (Quebec), section 12.1, added by Law 25 and in force since September 2023. It requires organizations to inform a person of such a decision and, on request, to explain the personal information used and the principal factors and parameters that led to it.

Builds on: *the CIA ERM resource's bias and fairness validation (2026); the AMF guideline's explanation requirement; established fairness-testing practice.*

Choosing a fairness measure is a judgment, not a calculation. The common measures answer different questions and, when underlying risk differs between groups, generally cannot all be satisfied at once.[7]

Table 3.5. Common fairness measures and their fit with actuarial practice

Measure	What it checks	Fit with actuarial practice
Demographic parity	Groups receive the same average outcome, such as the same average premium	Ignores genuine risk differences; generally incompatible with risk-based pricing
Equal error rates (equalized odds)	Groups have the same rates of false positives and false negatives	Well suited to yes/no decisions such as fraud flags, claim triage, or underwriting referrals
Calibration within groups	A prediction means the same thing in every group — a predicted 10% claim rate is 10% in each	Closest to actuarial fairness; compatible with different average prices that reflect risk
Disparate impact ratio	The ratio of favourable-outcome rates between groups	A screening test that signals where to look more closely, not a conclusion
Individual fairness	Similar risks are treated similarly	Intuitive, but depends on how "similar" is defined

For organizations: record which fairness measures were chosen for each model and why, and report results to the board for high and critical tier models.

For the CIA: respond to the bias consultation by October 15, 2026, and develop practical guidance on fairness testing for provincial rate filings and for underwriting and claims decisions.

Questions for leadership

- Should the CIA state a preferred fairness measure for insurance pricing, or leave the choice to the actuary with disclosure?
- How should actuaries test outcomes for groups whose characteristics insurers do not collect?

[7] This is a known result in the fairness literature: when underlying risk differs between groups, a model cannot in general be both calibrated within each group and have equal error rates across groups. See Kleinberg, Mullainathan, and Raghavan, "Inherent Trade-Offs in the Fair Determination of Risk Scores" (2016), and Chouldechova, "Fair Prediction with Disparate Impact" (2017).

3.4 Third parties, concentration, and silent AI

The issue. Most AI used by insurers and pension plans is bought, not built. That moves much of the risk outside the organization's direct control, and into a small number of providers.

What the evidence says

- **Concentration.** The CIA ERM resource warns that "reliance on a small number of global AI and cloud service providers increases the risk of a single point of failure," and that most of these providers "operate outside traditional regulatory frameworks."
- **Silent AI.** It also notes that "many vendors are embedding AI capabilities into products through automatic or mandatory rollouts, meaning the technology may be in use without sufficient organizational awareness, governance or testing."
- **Regulation.** OSFI's Guideline B-10 on third-party risk applies to AI vendors, and E-23 requires identifying and managing third-party models and data.

Looking ahead

Silent AI is likely to follow the path of silent cyber: exposure that no one priced or governed, discovered after a loss. A second risk is commercial — a provider changing its terms, prices or models in ways that leave the user unable to continue or to switch quickly.

Recommended approach

The CIA resource sets out the questions and contract protections organizations should seek. We have arranged them as a checklist for procurement and renewal.

Table 3.6. Checklist for AI in third-party services (based on the CIA ERM resource)

Area	What to require or confirm
Disclosure	Whether the product uses AI, which models and providers, and advance notice before AI features are added or changed
Data use	Whether our data is used to train the provider's models; where data is stored and processed
Performance	Testing results, known limitations, and performance metrics; the right to test independently
Audit and incidents	Audit rights, breach notification, access to logs
Change and exit	Notice of changes to terms, pricing, or models; data return and transition support; a manual fallback
Liability	Indemnities, performance warranties, and intellectual property terms

We also propose a regular **silent AI sweep** (proposed by this paper): each quarter, the owners of critical software review release notes and vendor notices for new AI features and add anything found to the AI inventory for scoring. It is a small, low-cost control for a risk that otherwise goes unseen.

Builds on: *the CIA ERM resource's warning about automatic vendor rollouts of AI (2026); the insurance market's experience with silent cyber.*

Questions for leadership

- Should actuaries report AI provider concentration explicitly in FCT and ORSA reports?

3.5 AI in capital, FCT and ORSA

The issue. Current guidance describes AI risks but does not measure them. Yet the actuarial tools for doing so — financial condition testing (FCT), the own risk and solvency assessment (ORSA) and stress testing — already exist.

What the evidence says

- The IFoA and London Foundation for Banking and Finance survey (May 2026) found 70% of senior financial services practitioners see AI risks as among the greatest facing their sector over the next five years.
- The CIA ERM resource identifies ways AI could affect each risk category: AI-amplified market volatility, AI making it easier to move deposits and pension transfers (reducing their "stickiness"), and credit losses as AI reshapes industries. It also warns that AI used to generate stress scenarios is itself a source of model risk.
- Neither the CIA resource nor the other guidance reviewed for this paper proposes specific AI scenarios for capital testing.

Looking ahead

Supervisors are likely to expect AI scenarios in stress testing within a few years, as happened with climate risk. Organizations that build them now will shape what good practice looks like.

Recommended approach

We propose that FCT and ORSA include at least two AI scenarios, chosen from the illustrative set below according to the organization's business. Severity will rest on judgment until experience accumulates, so the scenarios are also well suited to reverse stress testing: how bad would the event need to be to threaten solvency?

Table 3.7. Illustrative AI scenarios for FCT and ORSA (proposed by this paper)

Scenario	What happens	What it tests	Possible management actions
Provider failure	A critical AI or cloud provider fails, or a faulty model update spreads, for several weeks	Operational resilience; claims and underwriting backlogs; expenses	Manual fallback; second provider; communication plan
Shared-model error	An AI pricing or reserving model used across a line understates costs, and the error is found two to three years later	Rate adequacy; adverse reserve development; capital	Rate action; reserve strengthening; reinsurance
AI-assisted switching	Customers' AI agents compare and switch products at scale; lapse and transfer rates rise sharply	Lapse-sensitive liabilities; liquidity; recovery of acquisition costs	Product redesign; liquidity buffers; retention offers
Correlated market shock	AI-driven trading amplifies a market fall and a liquidity squeeze	Asset values; ALM; liquidity	Hedge review; liquidity buffers; asset mix
Fairness finding	A regulator or court finds an AI rating or claims model unfairly discriminatory	Redress and refunds; repricing; reputation	Model replacement; provisions; customer remediation

Builds on: *the CIA ERM resource's analysis of how AI affects each risk category (2026); existing FCT and ORSA scenario practice.*

Two further points. Validate AI-generated stress scenarios as models in their own right before use. And AI operational risk should be reflected in operational resilience planning under OSFI's Guideline E-21.

For the CIA: include AI scenario examples in future FCT guidance and educational notes.

Questions for leadership

- Should AI operational risk carry an explicit capital charge, or remain within existing operational risk capital?

3.6 Section 3 at a glance

Table 3.8. Summary of recommendations for organizations

Audience	Recommendations
Organizations	Keep one AI inventory scored on the CIA's five factors, with a severity-dominant rule. Govern acting AI as a digital delegate. Scale validation to each use case's risk tier. Test fairness on outcomes. Apply the third-party checklist and a quarterly silent AI sweep. Add at least two AI scenarios to FCT and ORSA. Prepare now for E-23 and the AMF guideline, both effective May 1, 2027.

Audience	Recommendations
Actuaries	Design or approve validation of AI affecting pricing, reserving, capital, claims, or benefits. Document the choice of fairness measures. Report AI concentration risk.
CIA and SOA	Publish a joint classification test. Develop guidance on validating AI, including agentic systems. Respond to the bias consultation by October 15, 2026, and issue fairness testing guidance. Add AI scenarios to FCT guidance.

4. AI as a risk and a market

Sections 2 and 3 treated AI as a tool that actuaries and their organizations use. This section turns to AI as something actuaries must measure in the world: a new source of insured losses, and a force that can shift the behaviour, markets, and assumptions on which actuarial models rest. Both sides also create new work for actuaries.

4.1 AI as an insurable peril

The issue. As businesses deploy AI, they take on new liabilities when it fails: wrong advice, unfair decisions, infringed rights and, increasingly, actions taken by AI agents on their behalf. Insurers must decide whether to exclude these losses, cover them silently or price them openly. All three choices are actuarial problems.

What the evidence says

- **Liability follows the deployer.** In *Moffatt v. Air Canada*, the British Columbia Civil Resolution Tribunal held Air Canada liable for incorrect fare information given by its website chatbot, rejecting the argument that the airline was not responsible for what the chatbot said. The award was small, but the principle is significant: an organization answers for its AI's statements as it does for its own.[8]
- **Exclusions are spreading.** In the United States, optional ISO endorsements allowing insurers to exclude generative-AI losses from commercial general liability policies became available in January 2026.[9]
- **Affirmative cover is emerging.** Munich Re has offered AI insurance since 2018; its aiSure product covers performance warranties given by AI vendors, AI liabilities such as discrimination, intellectual property infringement and hallucinations, and financial losses from AI errors. In April 2025, Armilla Insurance Services launched AI liability cover underwritten at Lloyd's by Chaucer, responding to AI that fails to

[8] Moffatt v. Air Canada, 2024 BCCRT 149 (February 14, 2024). The tribunal found that Air Canada did not take reasonable care to ensure its chatbot was accurate, and awarded $812.02 including interest and fees.
[9] Verisk/ISO endorsements CG 40 47 (Coverages A and B), CG 40 48 (Coverage B only), and CG 35 08 (products and completed operations), all optional at each insurer's discretion. Canadian policy wordings should be reviewed separately.

perform as intended or produces errors, hallucinations, or inaccuracies. In 2026, Corgi launched a dedicated AI insurance product (May), and CFC added explicit, affirmative AI language across seven policy lines, including technology and professional errors and omissions (June).

- **Underwriting is hard.** The Casualty Actuarial Society's *Actuarial Review* (May/June 2026) describes AI agents creating complex, difficult-to-underwrite exposures across cyber and professional lines.
- **No loss history.** Affirmative AI covers are new, so credible claims experience is still limited. Early pricing must rely on judgment, analogy with cyber and professional liability, and scenario loads.

Looking ahead

AI liability appears to be following the path cyber insurance took: losses first sit silently in existing policies, then insurers exclude them, then a specialist market grows to cover them affirmatively. Agentic AI will push the exposure into more lines — professional and directors' liability, product liability and, as vehicles become more autonomous, auto.

The largest unknown is **accumulation**. Most AI applications are built on a small number of foundation models from a few providers. A flaw in one widely used model, or a faulty update, could cause losses for thousands of insureds at once — the AI equivalent of a catastrophe event.

Recommended approach

The starting point is to separate AI losses into perils with different loss profiles, so each can be priced with the most suitable analogy.

Table 4.1. AI perils, loss profiles, and pricing analogues (proposed by this paper)

Peril	Example	Lines affected	Expected loss profile	Closest pricing analogue
Erroneous output	A chatbot misstates policy terms or fees	Professional liability; general liability; directors' liability	Frequent; small to medium; reporting delays of months to years	Professional and technology errors and omissions
Biased or unfair decisions	AI used in hiring, lending, pricing, or claims treats a group unfairly	Employment practices; directors' liability; regulatory	Infrequent but clustered; class actions; long tail	Employment practices and class action experience

Peril	Example	Lines affected	Expected loss profile	Closest pricing analogue
Intellectual property and privacy	Output infringes copyright; confidential data leaks through an AI tool	Media liability; cyber	Severity driven; uncertain law	Media and cyber liability
Autonomous action	An AI agent places wrong orders or trades; an automated vehicle causes a collision	Errors and omissions; product liability; auto	Low frequency; high severity	Product liability and auto
Performance shortfall	An AI model fails to meet contracted accuracy	Performance warranties; contractual liability	High frequency; low severity	Warranty and credit-style covers
Systemic model failure	A flawed update to a widely used model affects many users at once	All of the above, at the same time	Accumulation event	Cyber catastrophe scenarios

Builds on: *the CIA ERM resource (2026) on AI risk categories; published descriptions of Munich Re aiSure and Armilla/Chaucer AI cover; the development of the cyber insurance market.*

For pricing a peril with little or no claims history, we propose a staged approach that makes the basis of the rate transparent and moves weight to experience as it accrues (proposed by this paper):

- **Stage 1 — judgment and analogy.** Start from the closest analogue line, add explicit loads for AI-specific features, and document every assumption.
- **Stage 2 — exposure and scenarios.** Rate on measures of AI exposure (for example, the risk tier of the insured's AI uses) and test accumulation with scenarios, such as the failure of a major foundation model.
- **Stage 3 — credibility.** Blend in the line's own experience as it emerges, using standard credibility weighting: $Z = n / (n + k)$, where n is the volume of experience and k reflects how variable it is. Review the blend at least annually.

Builds on: *standard actuarial credibility theory; the cyber insurance market's move from analogy to experience-based pricing.*

Three further points follow:

- **Track accumulation by provider.** Record which foundation models and AI providers each insured depends on, and set aggregate limits by provider, as catastrophe exposures are managed by region.

- **Underwrite the insured's AI governance.** The CIA's five-factor scorecard (Table 3.1) and the severity-dominant rule (Table 3.2) can be used by underwriters to assess an insured's AI uses, rewarding good governance with better terms.
- **Make the wording clear.** Each policy should either cover AI losses affirmatively or exclude them clearly. Silence creates the unpriced exposure that troubled insurers in the early years of cyber.

There is also a Canadian opportunity. Where the commercial market for AI liability is thin, captive and association structures can let groups of businesses with similar AI exposures pool and retain risk while loss data builds, as captives did in the early days of other emerging covers.

Captives and AI risk (proposed by this paper)

Why captives fit. Commercial cover for AI liability is new, thin, and increasingly subject to exclusions. A captive lets an organization, or a group with similar exposures, fund and retain AI risk on its own terms, buy reinsurance for the tail, and build the loss data the commercial market lacks.

Possible structures. A single-parent captive for a large organization that deploys AI widely; a group or association captive for businesses with similar AI exposures, such as firms in one industry using comparable models; and, for measurable failures, parametric covers that pay when an AI system falls below an agreed performance level.

The Canadian setting. British Columbia has long had a captive regime, and Alberta's *Captive Insurance Companies Act* has been in force since July 1, 2022, including association captives. Both give Canadian organizations a domestic option.

The actuarial role. Feasibility and loss modelling with thin data; premium and reserve setting; capital and stress testing, including accumulation from shared foundation models; reinsurance design; and governance. A captive can make sound AI governance, such as the scorecard in Table 3.1, a condition of cover.

Caution. Members of a group captive may all rely on the same model or provider, so one failure can hit every member at once; aggregate limits and reinsurance are essential. Legal liability for AI is still developing, and regulatory and tax requirements must be met. A captive is a way to manage AI risk, not to avoid governing it.

Questions for leadership

- Should the CIA, SOA and CAS sponsor an industry database of AI incidents and claims with a common taxonomy, as the operational risk community did for operational losses?

- Should actuaries be expected to report "silent AI" exposure in their insurers' portfolios, as many now do for cyber?

4.2 Systemic and second-order effects

The issue. AI changes more than the losses insurers cover. It can change how customers behave, how markets move and how industries fare — and so change the assumptions behind lapse rates, claims patterns, asset returns and credit losses. These effects are easy to miss because they arrive gradually, through many small decisions.

What the evidence says

- **Financial Stability Board.** Its November 2024 report on the financial stability implications of AI identified vulnerabilities including reliance on a few dominant providers for AI hardware and cloud services, market correlations from "widespread use of similar AI models and training data," cyber risk and model risk. It also warned that generative AI could enable disinformation "capable of triggering flash crashes or bank runs."
- **CIA ERM resource.** It notes that AI "should make it easier to move deposits between financial institutions" and "could also facilitate transfers from pension plans, potentially affecting historical experience;" that AI-driven trading could amplify volatility through herding and feedback loops; that AI "has the potential to create winners and losers across a range of industries, resulting in credit losses;" and that AI's energy use may affect an organization's ability to meet its climate commitments.
- **International Energy Agency.** Its April 2025 report projects global data centre electricity demand to more than double by 2030, to around 945 terawatt-hours — slightly more than Japan's entire electricity use today.
- **IFoA and LFBF.** Their 2026 framework groups AI risks into outcomes, operating environment, and system risks, and argues that many of the hardest AI risks are system-level, scaling and spreading across the financial system.

Looking ahead

Customers will increasingly use AI agents to compare products, switch providers, file claims, and challenge decisions. Behaviour that actuaries have long treated as stable — inertia, persistency, and claim reporting patterns — may shift quickly, and in the same direction across many policyholders at once. On the asset side, many investors using similar AI models could make markets move together more often in stress.

Recommended approach

We recommend that actuaries review the assumptions most exposed to AI-driven change and act now, before the effects are visible in the data.

Table 4.2. Second-order effects of AI on actuarial assumptions (proposed by this paper)

Effect	How it happens	Assumptions affected	What to do now
AI-assisted switching	Customers' AI agents compare products and move money with little effort	Lapse, persistency, pension transfers, option take-up	Add an AI-switching sensitivity; track experience by sales and service channel
Changed claims behaviour	AI helps claimants file and contest claims; insurers' AI triage changes settlement patterns	Claim incidence, severity, reporting delays, development patterns	Test loss triangles for breaks when claims processes change; separate data before and after the change
Synthetic fraud	Deepfake documents, images, and voices support false claims	Fraud allowances, claim severity	Update fraud assumptions; invest in detection; share intelligence across the industry
Correlated markets	Many investors use similar models and data	Correlations, tail dependence, and liquidity in economic scenarios	Stress correlations upward in tail scenarios; hold liquidity buffers
Reshaped credit	AI creates winners and losers across industries	Default and migration rates by sector	Review credit portfolios by sector exposure to AI disruption
Energy and climate	Rapid growth in data centre electricity demand	Climate transition scenarios; infrastructure investments; property risk	Include AI-driven energy demand in climate scenarios
Health and longevity	AI may speed diagnosis and drug discovery	Mortality improvement; longevity risk	Add an upside longevity scenario; this effect is speculative and should be monitored, not assumed

Builds on: *the FSB report on AI and financial stability (2024); the CIA ERM resource (2026) on AI effects by risk category; the IEA's Energy and AI report (2025); the SOA's call for essays on AI and mortality.*

For organizations: add these effects to emerging-risk registers and to the sensitivity testing that supports valuations, FCT and ORSA.

For the CIA and SOA: add a field identifying AI-assisted channels to industry experience studies, so that changes in behaviour can be measured as they emerge.

Questions for leadership

- Which industry experience studies should start tracking AI-assisted customer behaviour now?
- Should the profession produce guidance on economic scenario generators in markets where many participants rely on similar AI models?

4.3 Where the opportunity lies

The risks in this section are also the profession's clearest growth area. Each needs the skills actuaries already have — pricing with limited data, reserving long-tailed liabilities, modelling accumulation, and setting assumptions under uncertainty:

- **Pricing and reserving AI covers**, from performance warranties to AI liability and its long-tailed claims.
- **Accumulation modelling** for foundation-model and provider concentration, for insurers, reinsurers, and catastrophe bond investors.
- **Underwriting assessment** of insureds' AI governance, using the scorecard approach in Section 3.
- **Captive and alternative risk transfer** for AI exposures the conventional market will not yet take.
- **Assumption setting and experience analysis** for AI-driven changes in customer behaviour, claims, and markets.
- **Expert evidence** in disputes over AI-related losses, where courts will need independent, quantitative opinions.

4.4 Section 4 at a glance

Table 4.3. Summary of recommendations on AI as a risk and a market

Audience	Recommendations
Insurers and reinsurers	Separate AI losses into perils and price each from its closest analogue. Track accumulation by AI provider. Underwrite insureds' AI governance. Make policy wording affirmative or excluded, never silent.
Actuaries	Use a staged, transparent approach to pricing AI perils. Review assumptions exposed to AI-driven change, starting with lapse, claims patterns, correlations, and credit. Add AI effects to sensitivity testing.
CIA, SOA and CAS	Sponsor an industry AI incident and claims database. Add AI-channel fields to experience studies. Develop guidance on economic scenarios where AI drives market correlation.

5. AI by practice area

Sections 2 to 4 set out themes that apply across the profession. Actuaries, however, work and qualify by practice area, and AI's effect differs sharply between them. This section translates the themes into what changes in each area, so that practitioners, employers, and education committees can see what matters most for their own work.

Each area follows the same short format: where AI is used now, what is likely to change by 2035, the risk most distinctive to that area, what actuaries should do, and the implications for the syllabus. The areas broadly follow the CIA's Fellowship specialty areas and the corresponding SOA and CAS pathways.

5.1 Where the themes bite hardest

The themes in Sections 2 to 4 do not weigh equally in every area. The table below gives our assessment of where each matters most. It is a judgment for discussion, not a measurement, and is intended to help leadership set priorities for guidance and education.

Table 5.1. Relative importance of each theme by practice area (proposed by this paper)

Practice area	Reliance (S.2)	Fairness (3.3)	Governance (3.2)	Behaviour (4.2)	Correlation (3.4, 4.2)	AI peril (4.1)
Life and annuities	High	High	High	Medium	Medium	Low
Group benefits and health	Medium	High	High	High	Low	Low
Pensions and retirement	High	Medium	Medium	High	Medium	Low
Property and casualty	High	High	High	Medium	Medium	High
Investment and ALM	Medium	Low	High	Medium	High	Low
ERM	High	Medium	High	Medium	High	High
Climate	Medium	Medium	High	Low	High	Medium

Builds on: *the themes and evidence in Sections 2 to 4 of this paper; the author's practice experience. Columns: reliance = accountability and reliance on AI; fairness = fairness and bias; governance = governance and validation; behaviour = shifts in customer and member behaviour; correlation = concentration and correlation; AI peril = AI as an insurable peril. Shading: darker = higher importance.*

Three patterns stand out. Fairness bites hardest wherever decisions are made about individuals — life, health and P&C. Behaviour shift matters most where outcomes depend on what customers and members choose to do — group benefits and pensions. Concentration and correlation dominate wherever the risk is systemic — investment, ERM and climate.

5.2 Life insurance and annuities

Where AI is used now. Accelerated, automated underwriting; review of medical records and attending physician statements; fraud detection; and automation of IFRS 17 and LICAT reporting. The SOA's expert panel on life underwriting (July 2026) found that AI "is already producing value, but that value is uneven, case-specific, and heavily influenced by carrier maturity, data readiness, workflow design" and staff skill.

What changes by 2035. Underwriting becomes more continuous, drawing on electronic health records, prescription histories, and wearable devices. Mortality and lapse models use many more variables. Reinsurers, who see experience across many insurers, are likely to set many of the standards for acceptable AI underwriting.

The distinctive risk. Long-duration products lock in assumptions for decades. If an opaque underwriting or pricing model is wrong, the error may take 10 to 20 years to emerge in mortality and lapse experience, long after the business has been written. A further Canadian constraint applies: the federal Genetic Non-Discrimination Act, upheld by the Supreme Court of Canada in 2020, prevents insurers from requiring genetic test results, and AI models must not reintroduce genetic information through proxies.[10]

What actuaries should do?

- Treat AI underwriting decisions as an assumption-setting input: monitor actual against expected mortality and lapse by underwriting path, from the first year.
- Require explainable reasons for every automated decline or rating, both for customers and for the appointed actuary.
- Test for proxies of protected characteristics, including genetic and health conditions that cannot lawfully be used.
- Work with reinsurers on shared validation of AI underwriting tools.

Syllabus implications. Explainable underwriting models; monitoring new-business assumptions under automated underwriting; fairness testing where health information is involved; and the effect of AI on long-term mortality and lapse assumptions.

5.3 Group benefits, disability, and health

Where AI is used now. Claims triage, predicting disability duration and return to work, case management support, and detection of drug, provider, and claim fraud. The SOA's 2026 expert panel on AI in health found actuaries using AI across claims triage, utilization

[10] Genetic Non-Discrimination Act, S.C. 2017, c. 3; Reference re Genetic Non-Discrimination Act, 2020 SCC 17.

management, provider monitoring, and fraud detection, and concluded that "AI was consistently viewed as an enabler of actuarial work rather than replacement for professional judgment." In Australia, the Actuaries Institute has examined the use of generative AI in injury and disability schemes.

What changes by 2035. AI-driven case management increasingly decides which claims receive which interventions, and when. Claimants and their advisers also use AI to prepare and contest claims.

The distinctive risk. Experience changes because the **claims process** changes, not because the insured population has changed. Disability termination rates, claim durations, and settlement patterns are central to reserving, and they can shift quietly when AI changes how claims are managed. Reserving bases calibrated on past claims management may then no longer fit, and development-based methods may misread the change as random variation.

What actuaries should do

- Record when and how AI changes claims management, and treat those dates as potential breaks in experience.
- Monitor termination and recovery rates separately for claims managed with and without AI support.
- Use claim-level methods alongside development-based diagnostics for long-duration disability liabilities, so that changes in termination patterns are seen directly.[11]
- Test for bias in health data, which the SOA panel flagged as reflecting socioeconomic status, race, and gender.

Syllabus implications. reserving and experience analysis when the claims process changes; claim-level valuation methods for disability; fairness in health and claims decisions; and working with clinicians and data scientists.

5.4 Pensions and retirement

Where AI is used now. Automation of valuations and benefit calculations, data validation, member communication and investment analysis. Pension funds also use AI for fraud analytics.

[11] See Lalani, M. H. (2026), "Reserving for Long-Term Disability Claims in a Property and Casualty Insurance Environment: A Hybrid Framework Integrating Tabular Cash-Flow Valuation and Development-Based Diagnostics," CIA contributed research, October 2026.

What changes by 2035. Members increasingly use AI advisers to decide when to retire, whether to transfer out, and how to draw down their savings. Plan sponsors and administrators deploy AI tools to answer member questions and support decisions.

The distinctive risk. Assumptions about member behaviour — retirement timing, commuted value and transfer elections, and option take-up — may become less stable. The CIA's ERM resource notes that AI "could also facilitate transfers from pension plans, potentially affecting historical experience." There is also a governance gap: federally regulated pension plans are excluded from OSFI's Guideline E-23, so pension AI governance depends on plan administrators' own risk management and on guidance from pension regulators.

What actuaries should do?

- Add sensitivity testing for AI-assisted member decisions, particularly transfer and commutation elections.
- Track election rates by communication channel to detect shifts early.
- Advise plan administrators on the fiduciary implications of member-facing AI, including accuracy, fairness, and record-keeping.
- Encourage pension regulators to issue AI governance expectations for plan administrators.

Syllabus implications. Behavioural assumptions for members using AI; governance and fiduciary duties when administrators deploy AI; and validation of automated benefit calculations.

5.5 Property and casualty insurance

Where AI is used now. P&C is the most advanced area. Machine-learning pricing, telematics, image-based claims assessment, fraud detection, catastrophe modelling and, increasingly, machine-learning approaches to reserving are in use. The CAS's 2026 Ratemaking, Product and Modelling seminar included sessions on machine learning, generative and agentic AI, regulation, and the role of the actuary.

What changes by 2035. Pricing moves closer to real time. In auto, driver-assistance and autonomous systems shift liability from drivers toward manufacturers and software providers, changing both the exposure and who pays for it. AI liability emerges as a new line (Section 4.1).

The distinctive risk. Two risks stand out. First, rate filings built on complex models that provincial rate regulators cannot readily follow or approve. Second, loss triangles are distorted when AI changes claims handling, so past development patterns no longer apply.

What actuaries should do?

- Prepare explainable support for machine-learning rate filings: what each variable does, why it is justified and how fairness was tested.
- Test loss triangles for breaks when claims processes change, and separate data before and after the change.
- Build the pricing and accumulation capability needed for AI liability and for autonomous vehicle exposures.
- Apply the CIA scorecard (Table 3.1) to underwrite insureds' AI governance.

Syllabus implications. Filing and explaining machine-learning models to provincial regulators; reserving under changing claims processes; pricing AI liability; and the changing liability landscape for automated vehicles.

5.6 Investment and asset-liability management

Where AI is used now. Economic scenario generation, asset screening, credit assessment, hedging support, and investment research.

What changes by 2035. A growing share of market activity is driven by AI models trained on similar data. AI also reshapes industries, creating winners and losers in credit portfolios.

The distinctive risk. Correlated behaviour. The Financial Stability Board identifies market correlations from "widespread use of similar AI models and training data" as a vulnerability, and the CIA's ERM resource notes that AI-driven trading could amplify volatility through herding. AI-generated economic scenarios are models that need validation.

What actuaries should do

- Stress correlations and liquidity in tail scenarios to reflect AI-driven herding.
- Validate AI-generated economic scenarios before using them in ALM, FCT or ORSA.
- Review credit portfolios for sector exposure to AI disruption.

Syllabus implications. Validating economic scenario generators that use AI; correlation and liquidity stress testing; and credit analysis of AI-driven industry change.

5.7 Enterprise risk management

ERM is the integrating discipline for this paper. The CIA's ERM resource argues that AI is "a transverse risk that affects all traditional risk categories," and that "characterizing AI as merely an emerging risk is increasingly inadequate." The practical guidance is in earlier sections: classification and governance (3.1 and 3.2), third-party and concentration risk (3.4), AI scenarios for FCT and ORSA (3.5), and systemic effects (4.2).

Syllabus implications. AI as a cross-cutting risk driver; AI risk appetite and reporting to boards; AI scenarios in capital and solvency testing; and model risk management under OSFI E-23 and the AMF guideline.

5.8 Climate risk

Where AI is used now. AI improves physical risk modelling by downscaling climate projections to local levels, building catastrophe event sets, and analyzing satellite and sensor data.

What changes by 2035. AI becomes central to climate risk assessment. At the same time, AI's own energy demand grows: the International Energy Agency projects data centre electricity use to more than double by 2030.

The distinctive risk. Over-confidence in AI projections that extrapolate beyond any observed climate, and the tension between AI's energy footprint and organizations' climate commitments, which the CIA's ERM resource also notes.

What actuaries should do

- Validate AI climate models against physical understanding and historical events, not only fit to data.
- Disclose the uncertainty in AI-based climate projections clearly.
- Include AI-driven energy demand in transition scenarios.

Syllabus implications. Validating AI climate and catastrophe models; communicating deep uncertainty; and the interaction between AI and climate transition risk.

5.9 Other areas to watch

- **Cyber and AI liability.** This is emerging as a specialty in its own right, combining pricing with thin data, accumulation modelling, and close knowledge of technology (Section 4.1).

- **Public plans and social security.** Governments are using AI in eligibility and benefit decisions. The public-trust stakes are high, and actuaries advising public plans should apply the fairness and explainability approaches in Section 3.3.

5.10 Implications for the syllabus by area

The table below brings together the syllabus implications, with an indication of material that can be reduced as AI takes over routine work. It feeds into the education recommendations in Section 6.

Table 5.2. Syllabus changes by practice area (proposed by this paper)

Practice area	Add or strengthen	Reduce or reframe
Life and annuities	Explainable underwriting; monitoring assumptions under automated underwriting; fairness with health information	Manual calculation of standard reserves and values; keep as foundations, assessed without AI
Group benefits and health	Experience analysis after process change; claim-level disability valuation; bias in health data	Routine experience-study production
Pensions and retirement	Member behaviour under AI advice; fiduciary governance of member-facing AI	Routine valuation and benefit calculation mechanics
Property and casualty	Filing and explaining machine-learning models; reserving under process change; AI liability and automated vehicles	Mechanical triangle methods taught in isolation from their limits
Investment and ALM	Validating AI scenario generators; correlation and liquidity stress	Routine data preparation for asset models
ERM and climate	AI as a cross-cutting risk; AI scenarios in FCT and ORSA; validating AI climate models	Descriptive coverage of AI as an "emerging" risk only

Builds on: *the practice-area analysis in this section; the SOA expert panel's view that obsolete methods should be retired as tools change, as commutation functions were.*

5.11 Section 5 at a glance

Table 5.3. Summary of recommendations by practice area

Audience	Recommendations
Practising actuaries	Identify where AI has changed processes in your area and treat those points as possible breaks in experience. Monitor assumptions most exposed to AI in your area: new business mortality, disability terminations, member elections, loss development, and market correlations.
Employers	Record AI-driven process changes so actuaries can adjust experience analysis. Give appointed actuaries explainable information on AI used in underwriting, claims, and pricing.

Audience	Recommendations
CIA, SOA and CAS	Use Table 5.1 to set priorities for practice-area guidance. Build the syllabus changes in Table 5.2 into each specialty track. Encourage pension regulators to set AI governance expectations for plan administrators.

6. The profession: workforce, education, relevance, and standards

The previous sections looked at the individual actuary, the organization, the market, and each practice area. This section turns to the profession as a whole: how many actuaries will be needed and in what roles, how they should be educated, where the profession can extend its reach, and how its standards should evolve. These are decisions that only the CIA, the SOA and their sister bodies can make.

6.1 The workforce: where work shrinks and where it grows

The issue. AI is taking over routine analytical work faster than any earlier technology. The question is not only how many actuaries will be needed, but at which levels — and how new actuaries will build judgment if the routine work they once learned from is automated.

What the evidence says

- **Junior hiring is falling in firms that adopt AI.** A Harvard study using US résumé and job-posting data found that junior employment in firms adopting generative AI declined by about 9% after six quarters relative to other firms, while senior employment held up. The decline came mainly from reduced hiring, not from layoffs.[12]
- **Early-career workers are most exposed.** A Stanford study found a 13% relative decline in employment for workers aged 22 to 25 in occupations most exposed to AI, while experienced workers in the same occupations remained stable or grew. The declines were concentrated where AI automates work rather than supporting it.[13]
- **Overall demand for actuaries remains strong.** The US Bureau of Labor Statistics projects employment of actuaries to grow 9% from 2025 to 2035, much faster than the average for all occupations, with about 1,500 openings a year.[14]
- **Professional bodies differ in their reading.** The Actuaries Institute expects task automation, more competition from data graduates and a decline in some roles. The SOA's expert panel judged that AI had not yet automated enough entry-level work to threaten training. The IFoA's President notes that each earlier technology reduced the number of actuaries needed per task, yet the profession grew as new opportunities opened.

[12] Hosseini, S. M. and Lichtinger, G., "Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data," Harvard University working paper (first version August 2025; revised May 2026).
[13] Brynjolfsson, E., Chandar, B. and Chen, R., "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Stanford Digital Economy Lab working paper (August 2025).
[14] US Bureau of Labor Statistics, Occupational Outlook Handbook: Actuaries (accessed October 2026).

Looking ahead

Both readings can be right at once. Total demand for actuaries may hold or grow while the **workforce shifts: fewer entry-level roles** and more roles in supervision, validation, and judgment. The risk is a delayed one. A thinner entry level today means fewer experienced Fellows in 10 to 15 years, just as demand for senior judgment rises.

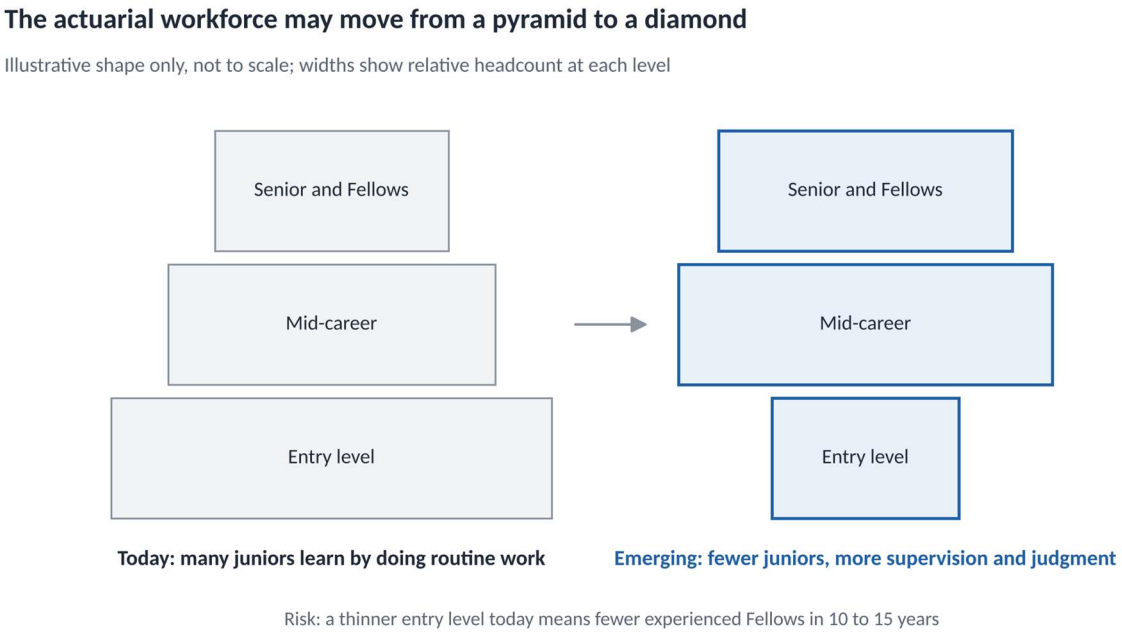


Figure 2. An illustrative change in the shape of the actuarial workforce.

Table 6.1. Where actuarial work is likely to shrink and to grow

Likely to shrink	Likely to grow
Data preparation, cleaning, and reconciliation	AI validation and model risk management
Routine valuation and reporting runs	Pricing and reserving AI liability and other new perils
First drafts of reports and documentation	Fairness and bias testing
Standard experience studies	AI scenarios, systemic risk, and accumulation modelling
Volume of entry-level analyst roles	Board, regulatory, assurance and expert-witness roles

Recommended approach

We propose that employers redesign entry-level roles around four elements, so that new actuaries still build judgment when AI does the routine work (proposed by this paper):

- **Foundations by hand.** Structured rotations in which core calculations are done from first principles before automation.
- **Supervising AI.** Early responsibility for checking AI output against the reviewability test in Section 2.1.
- **Challenge roles.** Assignments to test, break and validate models.
- **Early exposure to judgment.** Taking part in assumption setting, client discussions, and board materials sooner than is usual today.

Builds on: *the SOA expert panel's recommendations on apprenticeship (2026); the employment evidence above.*

We also propose a simple early-warning indicator: the CIA and SOA should track each year the ratio of new candidates entering actuarial employment to new Fellows qualifying, together with an employer survey of entry-level hiring plans. A sustained fall would signal pipeline stress years before it shows in the number of Fellows.

Questions for leadership

- If employers hire fewer juniors, who funds the apprenticeship that produces tomorrow's Fellows — employers, the profession, or the candidates themselves?
- Should the profession publish an annual workforce outlook, as some other professions do?

6.2 Education and qualification

The issue. Examinations that mainly test recall will lose meaning when every candidate has AI at work. Qualifications must certify what AI cannot do: judgment, supervision of AI, and accountability. It must also keep the foundations that make that judgment possible.

What the evidence says

- **Australia leads.** The Actuaries Institute makes Data Science Principles compulsory. Its Data Science Applications subject, covering neural networks and natural language processing, is chosen by over 60% of Fellowship students and is also offered as a microcredential to qualified members. All students also take Communication, Modelling and Professionalism.
- **The SOA has made AI a priority.** "Leverage AI" is one of three priorities in the SOA's 2025–2027 strategic plan, alongside the candidate pipeline and global growth.
- **Teaching must change with the tools.** The SOA expert panel warned that LLMs are "a very good productivity tool, but a very bad teaching tool," and recalled that

commutation functions remained on actuarial exams long after they ceased to be used: training should evolve as tools do.

- **Continuing education is already moving.** The US Qualification Standards require at least one hour of continuing education each year on bias, including bias in data, assumptions, algorithms, and models.

Looking ahead

Credentials will need to be refreshed through a career rather than earned once. The boundary between actuarial and data-science education will continue to blur, and the bodies that offer credible, modular AI education to both groups will shape the field.

Recommended approach

We propose an AI curriculum that runs from qualification through a career, tied to the competency ladder in Table 2.3 and the practice-area changes in Table 5.2.

Table 6.2. An AI curriculum by stage (proposed by this paper)

Stage	Content	Assessment
Foundation (before Associateship)	The main AI families and how each fails; data ethics and privacy; from GLMs to tree-based models and neural networks; core actuarial calculations from first principles	Without AI, to confirm understanding
Associateship	Applied machine learning in pricing and reserving; validation and explainability methods; using and checking LLMs; the reviewability test	With AI permitted; marks the candidate's critique of AI output
Fellowship (by specialty)	Practice-area content from Table 5.2; AI governance and model risk (OSFI E-23, the AMF guideline, the EU AI Act, the NIST AI Risk Management Framework); fairness testing; AI scenarios in FCT and ORSA; professionalism cases on reliance	Case-based, with AI permitted
After qualification	Supervising agentic AI; pricing AI perils; AI governance for boards; levels 3 and 4 of the competency ladder	Microcredentials and CPD

Builds on: *the Actuaries Institute's education program; the SOA expert panel (2026); the US Qualification Standards bias requirement; the competency ladder in Table 2.3.*

Three supporting changes follow:

- **A joint module.** The CIA and SOA should co-design one core AI module, recognized by both, rather than each building their own.

- **An AI and bias CPD requirement.** We propose at least one hour a year on AI and bias for practising members, modelled on the US requirement (proposed by this paper).
- **Retire what AI now does reliably.** Reduce content that tests routine mechanics, while keeping foundations assessed without AI (Table 5.2).

Questions for leadership

- Should candidates be allowed to use AI in some examinations, with marks for how well they challenge its output?
- Should the profession offer AI microcredentials to non-actuaries, as the Actuaries Institute is exploring?

6.3 Relevance and reach: the assurance profession for AI

The issue. If actuaries define their value by the calculations they perform, AI will erode it. If they define it by judgment, accountability and the public interest, AI expands it. The profession needs to choose, and say so publicly.

What the evidence says

- **Australia is extending its reach.** The Actuaries Institute's strategy promotes a "Data Science/AI Actuary" brand and is exploring whether to offer accreditation and support to the wider data professions. About one-third of its members specialize in data science.
- **The IFoA sees AI assurance as an actuarial role.** Its President argued that "AI is not going to replace actuaries, as it is actuaries that can make sure that AI is working as it should."
- **Actuaries must also learn from others.** The SOA expert panel observed that actuaries "don't have all the skills we need in a future state" and must act as translators between actuarial science and data science.

Looking ahead

Society will need trusted, regulated professionals to assure AI systems that make financial decisions about people. Auditors, engineers, and data scientists will compete for that role. The actuarial profession combines technical depth, enforceable professional standards, and a duty to the public interest — a combination none of the others fully shares.

Recommended approach

We recommend that the profession position the actuary as the **assurance professional for AI in financial decisions** and develop the roles that follow.

Table 6.3. Roles actuaries can own in an AI economy (proposed by this paper)

Role	What actuaries bring	Who else competes
AI validation actuary	Model risk expertise; understanding of pricing, reserving and capital consequences	Data scientists; model risk specialists
AI assurance and attestation	Independent, standards-based opinions, as appointed actuaries give today	Auditors; technology assurance firms
AI peril pricing and accumulation	Pricing with thin data; catastrophe-style accumulation modelling	Catastrophe modellers; insurtech underwriters
AI risk officer	Enterprise risk management; capital and scenario testing	Chief information security officers; compliance
Expert witness on AI losses	Quantification of loss and causation under uncertainty	Economists; forensic accountants
Public policy adviser	Evidence-based analysis of fairness, solvency, and social insurance	Policy analysts; academics

Builds on: *the Actuaries Institute strategy (2025–2027); the IFoA Presidential Address (2025); the themes in Sections 3 and 4 of this paper.*

In particular, we propose that the profession develop an **actuarial AI attestation**: an independent, standards-based opinion on whether a high-impact AI system used in insurance or pensions is sound, fair, and well governed. It would be modelled on the appointed actuary's opinion and could meet growing regulatory demand from OSFI, the AMF, and others (as proposed in this paper).

Questions for leadership

- Should the CIA, like the Actuaries Institute, consider offering credentials and standards to non-actuary AI practitioners?
- Should the profession seek recognition from regulators for actuarial AI attestations?

6.4 Standards and professional regulation

The issue. AI changes faster than standards can be written through full due process. Standards that are too specific will date quickly; standards that are too general will not guide practice.

What the evidence says

- **CIA standards.** In 2026, the CIA opened a review of the Standards of Practice for AI (Notice of Intent, May 26, 2026) and a review of bias in the General Standards, with comments due October 15, 2026, and final standards effective no later than May 1, 2027.
- **CIA Rules of Professional Conduct.** The 2026 consultation focuses on clarity and structure, conflicts of interest and modernized terminology that reflects work in multidisciplinary teams. Its published FAQ does not refer to AI or technology.
- **International.** The US Actuarial Standards Board is revising several ASOPs, and US actuaries already face an annual bias CPD requirement. The IAA published foundational papers on AI governance, testing, and documentation in March 2026, with more expected. The AAE has engaged directly with the EU on the AI Act.

Looking ahead

Regulators in Canada will apply OSFI E-23 and the AMF guideline from May 1, 2027. Professional standards should be ready to show how actuaries meet their part of those expectations, so that the profession leads rather than follows.

Recommended approach

We recommend addressing AI in layers, each changing at its own pace, with stable principles at the top and fast-moving guidance below.

Table 6.4. A layered approach to AI in professional standards (proposed by this paper)

Layer	Pace of change	AI content
Rules of Professional Conduct	Slow; changes rarely	Add AI-related commentary to the guidance accompanying the Rules: qualification to use AI, confidentiality of data entered into AI tools, and control of AI-drafted work
Standards of Practice	Multi-year due process	Technology-neutral principles on reliance, data, models, validation, and disclosure, as the Notice of Intent proposes
Educational notes	One to two years	Proportionate reliance on AI; fairness testing; AI validation; AI scenarios in FCT and ORSA
Practice resources and research	Months	Practice resources such as the CIA ERM resource; case studies; this paper
Qualification and CPD	Annual	AI curriculum (Table 6.2); AI and bias CPD; competency ladder (Table 2.3)

Builds on: *the CIA's 2026 reviews of its Standards of Practice and Rules of Professional Conduct; the structure of existing CIA guidance; the IAA's AI framework papers.*

Two further steps follow. The CIA should add AI-specific commentary to its draft guidance on the Rules of Professional Conduct before the 2026 consultation closes. The CIA, the US Actuarial Standards Board, and the IAA should align their AI wording so actuaries working across borders face one set of expectations.

Questions for leadership

- Can standard setting move at the speed of AI without losing due process, and is the layered approach the right way to achieve it?

6.5 Section 6 at a glance

Table 6.5. Summary of recommendations for the profession

Audience	Recommendations
CIA and SOA	Track the candidate-to-Fellow ratio and publish a workforce outlook. Co-design a joint AI module and adopt the AI curriculum by stage. Introduce an annual AI and bias CPD requirement. Develop an actuarial AI attestation and AI assurance roles. Address AI through layered standards, including AI commentary on the Rules of Professional Conduct.
Employers	Redesign entry-level roles around foundations, AI supervision, challenge roles, and early judgment. Keep hiring juniors and protect their training time.
Actuaries	Build AI skills through the competency ladder. Develop the roles in Table 6.3. Contribute to the 2026 consultations on standards and conduct.

7. Foresight: scenarios, the task survival model, and indicators to watch

No one can forecast how fast AI will develop or how society will govern it. Actuaries do not respond to that kind of uncertainty with a single prediction. Actuaries build scenarios, model the risks that matter, and watch for signals that show which path is unfolding. This section applies those methods to the profession itself.

This section has three parts: four scenarios for the profession to 2035; a task survival model that looks at actuarial work from an AI standpoint, task by task; and a set of early-warning indicators for leadership to monitor. It closes with principles for how actuaries can think and work in a world shaped by AI.

7.1 Four scenarios to 2035

The approach. Two uncertainties will shape the profession more than any others: how quickly AI capability advances, and how strong governance and public trust become. Crossing them gives four plausible futures. None is a forecast; together they mark out the range the profession should be ready for.

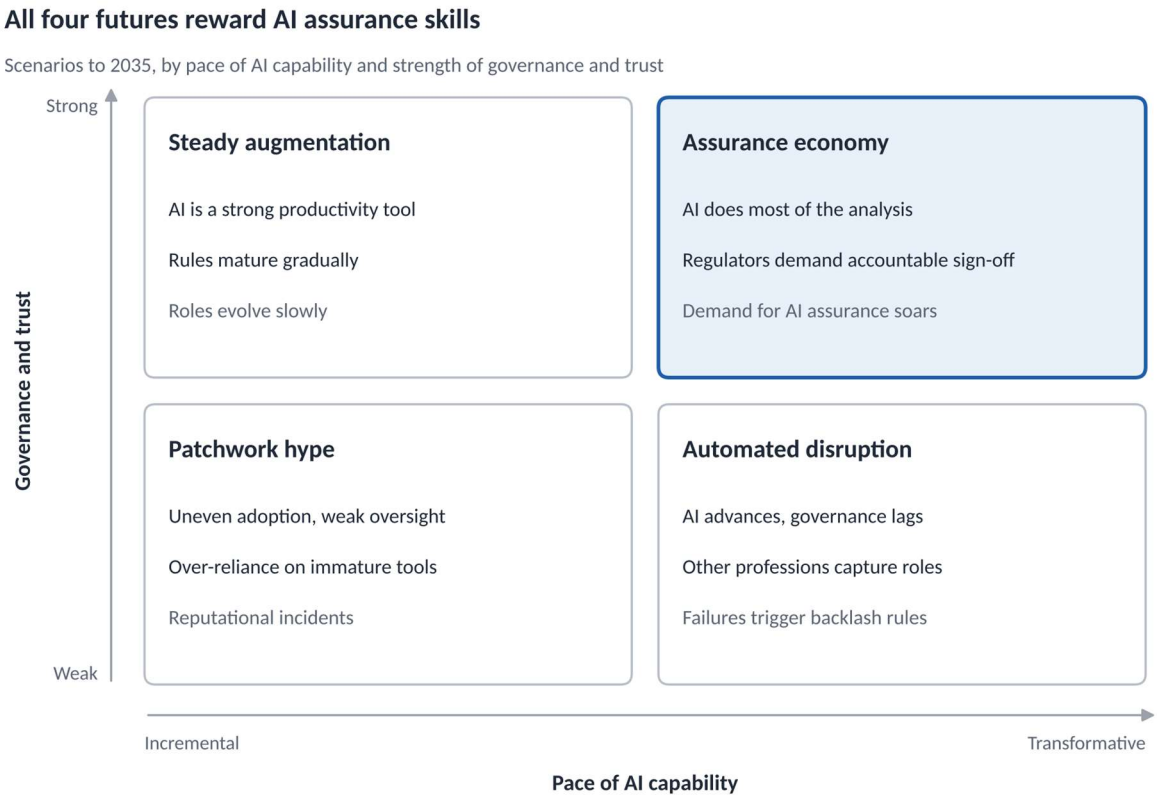


Figure 3. Four scenarios for the actuarial profession to 2035.

Table 7.1. The four scenarios and what they mean for actuaries (proposed by this paper)

Scenario	What it looks like	What it means for actuaries	Signposts that it is arriving
Steady augmentation	AI is a strong productivity tool; regulation matures gradually	Roles evolve slowly; efficiency gains; modest pressure on entry-level hiring	Adoption grows, but most AI remains advisory; regulators extend existing model-risk rules
Assurance economy	AI does most of the analysis; regulators require accountable human sign-off	Demand for validation, attestation and AI peril pricing rises sharply; fewer producers, more supervisors	Regulators ask for independent assurance of AI; affirmative AI insurance grows
Automated disruption	AI advances quickly; governance lags	Other professions and technology firms take actuarial roles; AI failures lead to abrupt, backlash regulation	Non-actuaries sign AI assurance work; major AI-driven losses or scandals

Scenario	What it looks like	What it means for actuaries	Signposts that it is arriving
Patchwork hype	Uneven adoption; weak oversight	Losses from over-reliance on immature tools; reputational damage to firms and professionals	Rising AI-related complaints and discipline cases; inconsistent rules across provinces and states

Builds on: *standard two-axis scenario planning; the evidence in Sections 1 to 6 of this paper.*

What the scenarios show. Three moves pay off in all four futures, and are therefore **no-regret** actions for the profession:

- **Build AI assurance skills** — validation, fairness testing, AI peril pricing, and attestation (Sections 3, 4 and 6).
- **Strengthen reliance and disclosure standards** — so that accountability holds however capable AI becomes (Sections 2 and 6.4).
- **Protect the apprenticeship** — so that the profession still produces experienced Fellows if entry-level work is automated (Section 6.1).

Beyond these, the profession can prepare **options** for specific scenarios without committing to them now: for example, a credential for non-actuary AI practitioners (valuable in the assurance economy), or a rapid-response standards process (valuable under automated disruption).

Questions for leadership

- Which scenario are the CIA and the SOA implicitly planning for today, and is that the right one?
- Which options should be prepared now so that they can be exercised quickly if the signposts appear?

7.2 The task survival model: actuarial work from an AI standpoint

The approach. Debates about whether AI will "replace actuaries" are unhelpful because AI does not replace occupations; it takes over tasks. We therefore propose looking at actuarial work as an AI system would — task by task — and applying a method every actuary knows: survival analysis.

Each actuarial task is treated as a life exposed to the "hazard" of automation. The hazard depends on four features of the task: how **codifiable** it is (C), how **data-rich** it is (D), how much **accountability** it carries (A) and how **novel** it is (N). The more codifiable and data-rich a task is, the sooner AI is likely to take it over; the more accountability and novelty it

carries, the longer it stays with people. In the proportional-hazards form (proposed by this paper):

$$\lambda_k(t) \;=\; \lambda_0(t)\exp(\beta_1 C_k + \beta_2 D_k - \beta_3 A_k - \beta_4 N_k)$$

$$S_k(t) \;=\; \exp\!\left(-\int_0^t \lambda_k(u)\,du\right) \qquad\qquad W(t) \;=\; \sum_k w_k\, S_k(t)$$

Here $\lambda_0(t)$ is the baseline pace of AI capability, which differs by scenario; C_k, D_k, A_k and N_k are scores for task k; $S_k(t)$ is the probability that task k still needs a person at time t; and W(t) is the share of today's actuarial work, weighted by the time spent on each task (w_k), that still needs an actuary.

An illustration

To show how the model works, we scored eight common tasks and applied two baseline paths: one for steady augmentation and one for the faster assurance economy. **The scores, time weights and parameters are assumptions chosen for illustration, not estimates.** The results show the shape of the answer, not its size.

Table 7.2. Illustrative task survival results for 2035 (proposed by this paper)

Task	C	D	A	N	Time weight	Still needs a person in 2035: Scenario: Steady augmentation	Still needs a person in 2035: Scenario: Assurance economy
Data preparation and reconciliation	4	4	2	2	20%	48%	11%
Routine valuation runs	4	4	3	2	15%	61%	22%
Experience studies	3	4	3	2	10%	72%	37%
Pricing model build	3	4	3	3	10%	80%	51%
Assumption setting	2	3	4	3	10%	93%	82%
Signed opinions and certification	2	3	5	3	10%	96%	87%
New-risk pricing (AI, climate)	2	2	4	4	10%	97%	91%
Board, regulator and client advice	2	2	4	4	15%	97%	91%
All tasks (weighted)					100%	77%	54%

Builds on: *illustrative assumptions only. C = codifiable, D = data-rich, A = accountability, N = novelty, each scored from 1 (very low) to 5 (very high); β = 0.4 for each feature; cumulative baseline hazard to 2035 of 0.15 (steady augmentation) and 0.45 (assurance economy). Results are not forecasts.*

Even with assumed inputs, the pattern is clear. Data preparation and routine valuation are largely automated in both scenarios. Assumption setting, signed opinions, new-risk pricing and advice to boards and regulators remain overwhelmingly human. The profession's total workload shrinks less than its routine workload, and the remaining work carries accountability and judgment.

How the numbers are calculated: a worked example

The calculation has five simple steps. We show them for one task — data preparation and reconciliation — under steady augmentation to 2035, and then add up all the tasks to get W.

Worked example: data preparation and reconciliation, steady augmentation, 2035

1. **Net task score.** C = 4, D = 4, A = 2, N = 2. Net score = C + D − A − N = 4 + 4 − 2 − 2 = 4.
2. **Relative speed of automation.** With β = 0.4 for each feature: exp (0.4 × 4) = exp(1.6) = 4.95. This task is automated about five times faster than a task with a net score of zero.
3. **Cumulative hazard to 2035.** We assume the baseline pace of automation under steady augmentation is 0.015 a year from 2025. Over 10 years, the cumulative baseline hazard is 0.015 × 10 = 0.15. For this task: 0.15 × 4.95 = 0.74.
4. **Survival.** S = exp (−0.74) = 0.48. About 48% of today's data preparation work still needs a person in 2035.
5. **Weighting.** Data preparation takes 20% of actuarial time today, so it contributes 20% × 0.48 = 9.5 percentage points to W.

Because the baseline pace is assumed constant here, the integral in the formula reduces to simple multiplication: annual baseline rate × years elapsed × relative speed. A more realistic model would let the baseline pace rise over time; the steps stay the same.

Where the time weights come from. The weights w_k are the shares of actuaries' working hours spent on each task today, and they must add to 100%. In this illustration, we have assumed the time weights. In practice, they would come from the member time-use survey recommended below.

Adding up. Repeating the five steps for every task and adding the contributions gives W, the share of today's work that still needs an actuary:

Table 7.3. How W is built up: steady augmentation, 2035 (proposed by this paper)

Task	Time weight (w)	Relative speed	Hazard to 2035	Still needs a person (S)	Contribution w × S (percentage points)
Data preparation and reconciliation	20%	4.95	0.74	48%	9.5
Routine valuation runs	15%	3.32	0.50	61%	9.1
Experience studies	10%	2.23	0.33	72%	7.2
Pricing model build	10%	1.49	0.22	80%	8.0
Assumption setting	10%	0.45	0.07	93%	9.3
Signed opinions and certification	10%	0.30	0.04	96%	9.6
New-risk pricing (AI, climate)	10%	0.20	0.03	97%	9.7
Board, regulator and client advice	15%	0.20	0.03	97%	14.5
Total (W)	100%				76.9

***Builds on** the illustrative assumptions in Table 7.2. Figures may not add exactly because of rounding.*

Over time. You can run the same calculation for any year by changing the years elapsed. Under the faster assurance economy, we assume a baseline pace three times higher, 0.045 a year:

Table 7.4. Share of today's work still needing an actuary (W), by year (proposed by this paper)

Year	Steady augmentation: cumulative baseline hazard	Steady augmentation: W	Assurance economy: cumulative baseline hazard	Assurance economy: W
2028	0.045	92%	0.135	79%
2030	0.075	87%	0.225	69%
2035	0.150	77%	0.450	54%

***Builds on:** illustrative assumptions only: baseline pace of 0.015 a year (steady augmentation) and 0.045 a year (assurance economy) from 2025; task scores, weights and β as in Table 7.2.*

How to read W. A W of 77% in 2035 means that about 23% of the hours actuaries spend today on these tasks would by then be done by AI. It does not mean 23% fewer actuaries. The new work described in Sections 4 and 6 — AI validation, assurance, AI peril pricing, and advice — is not in today's task list, and it adds to demand.

Recommended approach

We recommend that the CIA and the SOA turn this illustration into a real estimate (proposed by this paper):

- **Measure time use.** Survey members on how their time is divided across tasks, by practice area and career stage, to set the weights w_k.
- **Score tasks by expert panel.** Use structured expert judgment to score each task on the four features, and review the scores every two years.
- **Link the baseline to the scenarios.** Calibrate $\lambda_0(t)$ for each scenario, informed by the indicators in Section 7.3.
- **Validate against what happens.** Compare the model with hiring data, role changes, and the candidate-to-Fellow ratio (Section 6.1), and refine it.

The results would tell leadership, with numbers rather than opinions, where to shrink the syllabus, where to invest in new skills, and how quickly.

Questions for leadership

- Should the CIA and SOA commission a joint time-use survey of members as the first step?

7.3 Early-warning indicators

The approach. Scenarios are useful only if someone watches for signs that one is unfolding. We propose a short set of indicators, each linked to the scenario it points toward, for CIA and SOA leadership to review each year.

Table 7.5. Early-warning indicators for the profession (proposed by this paper)

Indicator	What a rise signals	Where to watch	Points toward
New candidates entering actuarial employment relative to new Fellows	Stress in the training pipeline	CIA and SOA exam and membership data; employer surveys	Automated disruption
Regulators using AI to review filings, or requiring independent AI assurance	The machine-review and assurance era is arriving	OSFI, AMF, provincial regulators, NAIC	Assurance economy
Affirmative AI insurance compared with AI exclusions	An AI peril market is forming	Policy wordings; carrier and reinsurer products	Assurance economy
Non-actuaries signing AI assurance work in insurance and pensions	Other professions are taking the role	Audit and consulting markets; regulatory filings	Automated disruption

Indicator	What a rise signals	Where to watch	Points toward
AI-related complaints and discipline cases involving actuaries	Over-reliance and weak review	CIA and US discipline summaries	Patchwork hype
Sudden, correlated shifts in lapse, transfer or claims experience	AI-driven customer behaviour	Experience studies; company monitoring	Any scenario with fast AI capability
Share of AI uses that act without human approval	Rising autonomy, and so rising governance demands	Organizations' AI inventories; industry surveys	Assurance economy or automated disruption
Concentration of AI use on a few foundation models	Rising systemic and accumulation risk	Vendor disclosures; insurer exposure data	Any scenario

Builds on: *the themes in Sections 3 to 6; the scenarios in Section 7.1.*

We recommend that the CIA and the SOA maintain these indicators as a simple **AI foresight dashboard**, reviewed annually by their boards alongside the scenarios, and used to adjust priorities for education, standards, and research (as proposed in this paper).

7.4 Thinking and working in an AI world

Foresight is also personal. Five principles can help individual actuaries think and work well as AI reshapes their field:

- **Think in scenarios, not forecasts.** Plan for a range of futures, and prepare for the ones that would hurt most.
- **Prefer options to bets.** Build skills and relationships that keep several career paths open, rather than staking everything on one.
- **Keep learning, continuously.** Treat AI capability as a moving target; refresh skills every year, not once a decade.
- **Guard the human core.** Judgment, ethics, accountability, and clear communication are what AI cannot supply, and what clients, boards and regulators will pay for.
- **Be the translator.** Stand between technical AI and the people affected by its decisions, and explain each to the other.

These principles are not new to the profession. Actuaries have always worked with uncertainty, long time horizons, and incomplete data. The AI world asks them to apply the same discipline to their own careers and to the future of their profession.

7.5 Section 7 at a glance

Table 7.6. Summary of recommendations on foresight

Audience	Recommendations
CIA and SOA	Act now on the three no-regret moves: AI assurance skills, stronger reliance and disclosure standards, and a protected apprenticeship. Prepare options for specific scenarios. Commission a joint time-use survey and calibrate the task survival model. Maintain an AI foresight dashboard that boards review annually.
Employers	Use the scenarios in workforce and capability planning. Share data on hiring and role changes to help calibrate the model.
Actuaries	Apply the five principles to your own career. Build skills for the work that remains human: judgment, accountability, new risks, and advice.

8. Conclusion

Every generation of actuaries has absorbed a new technology. Mainframes, personal computers, and predictive analytics each removed work that actuaries once did by hand, and each time the profession grew by moving to the new work the technology created. AI differs in one respect: for the first time, the technology reaches into drafting, reasoning, and decision-making—the core of professional work.

This paper has argued that the profession's response must work at three levels at once. **Individual actuaries** must remain accountable for AI-assisted work and build the judgment to supervise it. **Organizations** must govern AI in proportion to its risk, test it for fairness, and hold capital against it. **The profession** must reshape education, protect the apprenticeship, extend its reach into AI assurance, and keep its standards current. Alongside these, AI must be measured as a new source of insured, systemic, and behavioural risk.

Much remains uncertain: how fast AI will advance, how it will be governed, and how customers, markets and employers will respond. The scenarios in Section 7 show that the right response does not depend on resolving that uncertainty. The same three moves — building AI assurance skills, strengthening reliance and disclosure standards, and protecting the apprenticeship — pay off in every future we examined.

Table 8.1. A roadmap for the profession (proposed by this paper)

Horizon	Priority actions
Now to April 2027	Respond to the bias consultation and the Rules of Professional Conduct consultation. Begin the educational note on reliance on AI. Publish the joint classification test. Commission the time-use survey.

Horizon	Priority actions
By May 1, 2027	Organizations complete their AI inventories and governance for OSFI E-23 and the AMF guideline, with actuaries designing validation for AI in pricing, reserving, capital, and claims.
2027 to 2028	Launch the joint AI module, the AI and bias CPD requirement and the competency ladder. Pilot the actuarial AI attestation. Add AI scenarios to FCT guidance. Publish the first AI foresight dashboard and workforce outlook.
2028 to 2030	Embed the AI curriculum across all specialty tracks. Calibrate and update the task survival model every two years. Review standards against the emerging scenario.

Builds on: *the recommendations in Sections 2 to 7 of this paper.*

The title of this paper is deliberate. Artificial intelligence will become more capable, but it will not carry professional accountability, weigh competing interests or answer to the public. That is **actuarial intelligence**: judgment exercised under uncertainty, in the public interest, by people who stand behind their work. If the profession invests in it now, the AI world will need actuaries more, not less.

Appendix A. Glossary of AI terms

This glossary defines the AI terms used in this paper in plain language. Where the CIA's practice resource *Considerations for AI within an ERM Program* (September 2026) defines the same term, our wording is consistent. Terms are listed alphabetically.

Term	Definition (as used in this paper)
Agentic AI (AI agent)	AI that plans and carries out multi-step tasks toward a goal, using tools and taking actions (for example, updating a record or setting a price) with limited real-time human oversight.
AI lifecycle	The full life of an AI application: problem framing, data, development, validation, approval, deployment, monitoring, and retirement. OSFI's Guideline E-23 sets expectations at each stage.
AI system	A machine-based system that, with some degree of autonomy, infers from its inputs how to generate outputs such as predictions, content, recommendations, or decisions. See Section 1.3 and its footnote for how this differs from a "model" under OSFI E-23.
Automated decision-making	A decision about a person, such as an underwriting or claim decision, made by a system without meaningful human involvement.
Bias (algorithmic bias)	Systematic unfairness in a model's outcomes for some groups. It can persist even when protected characteristics are removed, through proxy variables or biased historical decisions.
Concentration risk (AI)	Exposure that arises when many firms rely on the same few AI or cloud providers, so that a single failure or change affects many of them at once.

Term	Definition (as used in this paper)
Explainability	The ability to explain why a model produced a given output, in terms a user or regulator understands.
Fairness metric	A measure of whether outcomes differ across groups, such as equal approval rates, equal error rates, or calibration by group. Different metrics can conflict, so choosing one is a judgment that should be documented.
Foundation model	A large model trained on broad data that can be adapted to many tasks. Most generative AI tools are built on a small number of them. The EU AI Act calls these general-purpose AI models.
Generalized linear model (GLM)	The standard statistical model for much actuarial pricing. It is interpretable and generally treated as a traditional actuarial model rather than AI, although the boundary is debated.
Generative AI	AI that produces new content — text, code, summaries, images, or audio — in response to instructions.
Gradient boosting	A machine-learning method that builds many small decision trees in sequence, each correcting the errors of the last; XGBoost is a common implementation. Widely used in pricing and fraud models.
Hallucination	Output from a generative AI tool that is fluent and confident but false or fabricated, such as an invented citation or figure.
Human-in-the-loop	A design in which a person reviews, approves, or can override an AI system's output before it takes effect.
Large language model (LLM)	A generative AI model trained on very large amounts of text to predict and produce language. The basis of chatbots and AI writing assistants.
Machine learning (predictive AI)	Methods that learn patterns from data to estimate an outcome, such as a claim cost or a lapse probability, rather than following rules written by hand.
Model	Under OSFI's Guideline E-23, a broad term covering any method that processes input data to generate results, including traditional actuarial models and AI or machine-learning models.
Model drift	A decline in a model's performance over time as the data or environment moves away from the conditions it was trained on.
Model risk	The risk of loss from decisions based on models that are wrong, misused or misunderstood.
Neural network	A model made of layers of connected nodes that learns complex, non-linear relationships from data.
Proxy discrimination	Unfair differences in outcomes that arise through variables correlated with a protected characteristic — for example, postal code — even when that characteristic is not itself used.
Silent AI	AI in use without the organization's knowledge, for example, when a vendor adds AI features through a routine software update. A term used in this paper, by analogy with "silent cyber."

Term	Definition (as used in this paper)
Third-party (vendor) model	A model or AI tool built by an outside provider. Under OSFI's Guidelines B-10 and E-23, the user remains responsible for managing its risk.
Training data	The data a model learns from. Its quality, coverage and history largely determine the model's accuracy and fairness.
Validation	Independent testing that a model is fit for its intended purpose, covering accuracy, robustness, fairness and explainability as relevant to its risk.

Appendix B. Abbreviations

This appendix spells out the abbreviations used in this paper. Technical AI terms are defined in Appendix A.

Table B.1. Abbreviations

Abbreviation	Full name	Note
AAE	Actuarial Association of Europe	Represents national actuarial associations in Europe
Academy	American Academy of Actuaries	US professional body; its practice councils set annual priorities
AI	Artificial intelligence	See Section 1.3 and Appendix A
ALM	Asset-liability management	
AMF	Autorité des marchés financiers	Quebec's financial sector regulator
ASB	Actuarial Standards Board (United States)	Sets the US Actuarial Standards of Practice. Canada has a separate Actuarial Standards Board, which sets the CIA Standards of Practice.
ASOP	Actuarial Standard of Practice	US standards set by the ASB
BCCRT	British Columbia Civil Resolution Tribunal	Cited in Section 4.1
CAS	Casualty Actuarial Society	North American body for property and casualty actuaries
CEO	Chief executive officer	
CIA	Canadian Institute of Actuaries	National body for the actuarial profession in Canada
CPD	Continuing professional development	
CRO	Chief risk officer	
ERM	Enterprise risk management	

Abbreviation	Full name	Note
EU	European Union	
FAQ	Frequently asked questions	
FCT	Financial condition testing	Canadian requirement to test an insurer's future financial position under adverse scenarios
FSB	Financial Stability Board	International body that monitors the global financial system
GLM	Generalized linear model	See Appendix A
IAA	International Actuarial Association	Global association of national actuarial bodies
IEA	International Energy Agency	
IFoA	Institute and Faculty of Actuaries	Professional body for actuaries in the United Kingdom
IFRS	International Financial Reporting Standards	IFRS 17 sets the accounting for insurance contracts
ISO	Insurance Services Office	US provider of standard insurance policy forms, part of Verisk
LFBF	London Foundation for Banking and Finance	Co-author, with the IFoA, of It's Still Not Magic (2026)
LICAT	Life Insurance Capital Adequacy Test	OSFI's capital framework for federally regulated life insurers
LLM	Large language model	See Appendix A
NAIC	National Association of Insurance Commissioners	Organization of US state insurance regulators
NIST	National Institute of Standards and Technology	US agency; publisher of the AI Risk Management Framework
ORSA	Own risk and solvency assessment	An insurer's own assessment of its risks and capital needs; required for federally regulated insurers under OSFI Guideline E-19
OSFI	Office of the Superintendent of Financial Institutions	Canada's federal regulator of banks, insurers, and federally regulated pension plans
P&C	Property and casualty	General insurance
SCC	Supreme Court of Canada	Used in case citations
SOA	Society of Actuaries	North American actuarial body serving life, health, pension, and investment actuaries
US	United States	

Table B.2. Regulatory guidelines, legislation, and standards referred to

Reference	Title	Note
OSFI Guideline B-10	Third-Party Risk Management	Applies to AI vendors and other third-party providers
OSFI Guideline E-19	Own Risk and Solvency Assessment	
OSFI Guideline E-21	Operational Risk and Resilience	
OSFI Guideline E-23	Model Risk Management (2027)	Effective May 1, 2027; brings AI and machine-learning models into scope
AMF guideline	Guideline for the Use of Artificial Intelligence	Effective May 1, 2027, for financial institutions authorized in Quebec
EU AI Act	Regulation (EU) 2024/1689 on artificial intelligence	High-risk obligations for life and health insurance pricing apply from December 2, 2027
Law 25 (Quebec)	Act to modernize legislative provisions as regards the protection of personal information	Amended Quebec's private-sector privacy legislation, including rules on automated decisions
ASOP No. 41	Actuarial Communications (US)	Under revision
ASOP No. 56	Modelling (US)	
IFRS 17	Insurance Contracts	International accounting standard for insurance contracts
OSFI LICAT guideline	Life Insurance Capital Adequacy Test	
NIST AI RMF	Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023	Voluntary US framework for managing AI risk

Appendix C. References

The sources cited in this paper are listed below by jurisdiction. Web addresses were current as of October 2026. Where a primary document could not be accessed directly, the summary used is identified.

Canada

Canadian Institute of Actuaries (2026). Considerations for AI within an ERM Program. Practice resource document, September 2026. https://www.cia-ica.ca/publications/226093e/

Canadian Institute of Actuaries (2026). Review of the Standards of Practice with Respect to the Use of Artificial Intelligence in Actuarial Practice – Notice of Intent. May 26, 2026. https://www.cia-ica.ca/publications/226065e/

Canadian Institute of Actuaries (2026). Bias considerations within Standards of Practice, Part 1000 – Notice of Intent. As reported by Insurance Portal: https://insurance-portal.ca/article/actuaries-asked-to-comment-on-bias-considerations/

Canadian Institute of Actuaries (2026). FAQ – Rules of Professional Conduct 2026 consultation. https://www.cia-ica.ca/being-a-member/member-voting/amendments/faq-rules-of-professional-conduct-2026-consultation/

Lalani, M. H. (2026). Reserving for Long-Term Disability Claims in a Property and Casualty Insurance Environment: A Hybrid Framework Integrating Tabular Cash-Flow Valuation and Development-Based Diagnostics. CIA contributed research, October 2026. https://www.cia-ica.ca/publications/226100e/

Office of the Superintendent of Financial Institutions. Guideline E-23, Model Risk Management (2027), published September 11, 2025, effective May 1, 2027; Guideline B-10, Third-Party Risk Management; Guideline E-19, Own Risk and Solvency Assessment; Guideline E-21, Operational Risk and Resilience. https://www.osfi-bsif.gc.ca

Autorité des marchés financiers (2026). Guideline for the Use of Artificial Intelligence. April 7, 2026, effective May 1, 2027. Summary: https://www.blakes.com/insights/quebec-s-amf-issues-guideline-on-the-use-of-ai-for-financial-institutions/

Act respecting the protection of personal information in the private sector (Quebec), section 12.1, as amended by Law 25 (in force September 2023).

Genetic Non-Discrimination Act, S.C. 2017, c. 3; Reference re Genetic Non-Discrimination Act, 2020 SCC 17.

Moffatt v. Air Canada, 2024 BCCRT 149 (February 14, 2024).

Government of Alberta. Captive Insurance Companies Act (in force July 1, 2022). https://www.alberta.ca/information-for-captive-insurers

Canadian Underwriter (2026). “Actuaries may have to defend AI-generated decisions.” June 8, 2026. https://canadianunderwriter.ca/news/actuaries-may-have-to-defend-ai-generated-decisions/

United States

Society of Actuaries Research Institute (2026). Actuarial Mindsets for Leading in the AI Era: An Expert Panel Discussion. January 2026. https://www.soa.org/48f844/globalassets/assets/files/resources/research-report/2026/2026-01-ait184-act-mindset.pdf

Society of Actuaries Research Institute (2026). AI and Life Underwriting in Transition: Insights from an Expert Panel. July 2026. https://www.soa.org/resources/research-reports/2026/ai-life-uw-transition/

Society of Actuaries Research Institute (2026). AI in Healthcare and Health Insurance: An Expert Panel Discussion. February 2026. https://www.soa.org/globalassets/assets/files/resources/research-report/2026/2026-02-ait171-ai-health.pdf

Society of Actuaries. Artificial Intelligence Research Reports. https://www.soa.org/research/topics/artificial-intelligence-res-report-list/

Society of Actuaries. Strategic Plan 2025–2027. https://www.soa.org/programs/strategic-planning

Society of Actuaries and Casualty Actuarial Society (2026). 19th Annual Emerging Risks Survey; media coverage at https://www.soa.org/resources/research-reports/2026/soa-research-in-news/

Casualty Actuarial Society (2026). Actuarial Review, May/June 2026. https://digital.casact.org/issue/may-june-2026/

American Academy of Actuaries (2026). “Agile and Ready for What’s Next” (2026 practice council priorities). https://actuary.org/article/agile-and-ready-for-whats-next/

Actuarial Standards Board (2026). Third exposure draft of a proposed revision of ASOP No. 41, Actuarial Communications. March 2026. https://actuary.org/asb-approves-a-third-exposure-draft-of-a-proposed-revision-of-asop-no-41/

Actuarial Standards Board. ASOP No. 56, Modelling. https://www.actuarialstandardsboard.org

American Academy of Actuaries. US Qualification Standards, including the annual bias continuing education requirement. Summary: https://www.casact.org/professional-education/clone-continuing-education-policy/us-qualification-standards

National Association of Insurance Commissioners (2023). Big Data and Artificial Intelligence (H) Working Group summary. https://content.naic.org/sites/default/files/national_meeting/BDAI%20Summary.pdf

National Association of Insurance Commissioners (2023). Model Bulletin: Use of Artificial Intelligence Systems by Insurers. https://content.naic.org

National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.

US Bureau of Labor Statistics. Occupational Outlook Handbook: Actuaries. https://www.bls.gov/ooh/math/actuaries.htm

Verisk/ISO (2026). Generative AI exclusion endorsements CG 40 47, CG 40 48, and CG 35 08. Summary: https://www.independentagent.com/vu_resource/verisk-to-roll-out-new-general-liability-exclusions-for-generative-ai-exposures/

United Kingdom and Australia

Institute and Faculty of Actuaries and London Foundation for Banking and Finance (2026). It's Still Not Magic: Framing the Risks Facing Financial Services in the Gen AI Era. May 28, 2026. https://actuaries.org.uk/news/its-still-not-magic-framing-the-risks-facing-financial-services-in-the-gen-ai-era/

Sweeting, P. (2026). IFoA Presidential Address 2025. British Actuarial Journal. https://www.cambridge.org/core/product/DF135117FD19071AD4BB8332DC124D96/core-reader

Actuaries Institute (2025). Strategic Plan 2025–2027: Reach and Impact in the Age of AI. https://content.actuaries.asn.au/resources/resource-ce6yyqn64sx3-2093352434-60119

Actuaries Institute (2024). The AI-Powered Actuary: Amplifying Our Value. https://content.actuaries.asn.au/resources/resource-ce6yyqn64sx3-2093352434-60069

Europe and international

Actuarial Association of Europe (2025). Navigating Europe's AI Act: Insights for Actuaries and the Insurance Sector. Discussion paper, March 2025. https://actuary.eu/wp-content/uploads/2025/03/AAE-Discussion-Paper-Navigating-Europes-AI-Act.pdf

Actuarial Association of Europe (2026). Press release supporting clarification of the classification of GLMs under the EU AI Act. April 15, 2026. https://actuary.eu/wp-content/uploads/2026/04/260415_EIOPA_GLM_Endorsement_PressRelease_Final.pdf

European Pensions (2026). "AAE signals readiness to support AI Act implementation." https://europeanpensions.net/ep/AAE-signals-readiness-to-support-AI-Act-implementation.php

Regulation (EU) 2024/1689 (Artificial Intelligence Act), as amended by Regulation (EU) 2026/1744 (Digital Omnibus on AI). Summary: https://cdp.cooley.com/digital-ai-omnibus-delays-key-deadlines-introduces-new-rules/

International Actuarial Association (2026). In Pursuit of Responsible AI: An Actuarial Profession Take (papers on governance, testing, and documentation). March 2026. https://actuaries.org/resources-post/in-pursuit-of-responsible-ai-an-actuarial-profession-take/

Financial Stability Board (2024). The Financial Stability Implications of Artificial Intelligence. November 14, 2024. https://www.fsb.org/uploads/R14112024.pdf

International Energy Agency (2025). Energy and AI. April 10, 2025. https://www.iea.org/reports/energy-and-ai

Research, market and other sources

Brynjolfsson, E., Chandar, B. and Chen, R. (2025). Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Stanford Digital Economy Lab. https://siepr.stanford.edu/publications/working-paper/canaries-coal-mine-six-facts-about-recent-employment-effects-artificial

Hosseini, S. M. and Lichtinger, G. (2026). Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data. Harvard University working paper, revised May 2026.

Chouldechova, A. (2017). Fair Prediction with Disparate Impact. Big Data, 5(2).

Kleinberg, J., Mullainathan, S., and Raghavan, M. (2016). Inherent Trade-Offs in the Fair Determination of Risk Scores. arXiv:1609.05807.

Parasuraman, R. and Manzey, D. H. (2010). Complacency and Bias in Human Use of Automation: An Attentional Integration. Human Factors, 52(3), 381–410.

KPMG (2026). 2025 Insurance CEO Outlook. January 2026. https://assets.kpmg.com/content/dam/kpmgsites/xx/pdf/2026/01/ceo-oulook-survey-insurance-report.pdf.coredownload.inline.pdf

Munich Re. aiSure: insurance for AI. https://www.munichre.com/en/solutions/for-industry-clients/insure-ai.html

Armilla (2025). Armilla launches affirmative AI liability insurance with Lloyd's underwriter Chaucer. April 30, 2025. https://www.armilla.ai/resources/armilla-launches-affirmative-ai-liability-insurance-with-lloyds-underwriter-chaucer

CFC (2026). CFC responds to customer demand for affirmative AI cover. June 25, 2026. https://www.cfc.com/en-au/knowledge/news/2026/06/cfc-responds-to-customer-demand-for-affirmative-ai-cover/

Corgi (2026). AI insurance launch, May 5, 2026, as reported by Glitchwire. https://glitchwire.com/news/corgi-launches-ai-insurance-as-model-risk-becomes-a-balance-sheet-problem/

Disclaimer and Disclosure of AI Use

The views, proposals and opinions in this paper are those of the author. They do not represent the views of the Canadian Institute of Actuaries, the Society of Actuaries or any other organization. Tools and frameworks marked “proposed by this paper” are offered for discussion. They have not been adopted or endorsed by any professional body.

In preparing this paper, the author used generative AI tools to search for and summarize published research, guidance and regulatory material, and to help draft text, tables and figures. The author reviewed all sources against the original documents. The author checked figures and references, and every finding, proposal, and conclusion reflects the author’s own professional judgment. The author takes full responsibility for the content.

This approach is consistent with the principles set out in this paper: AI can extend what an actuary can do, but accountability stays with the person behind the work.

This paper is a discussion document. It is not actuarial, legal or professional advice and should not be relied on as such for any particular situation. Information is current as of October 2026. Standards, guidance and legislation referred to here may change.