AI Risks in the Actuarial Profession

Insights from Senior Leaders in Phase 1 of the Foresight Project

August 2026

Authors

Kun Zhao

Alexander Saeri

Peter Slattery

Neil Thompson

Contents

Executive summary 3
1. Introduction 5
2. What influences AI adoption in the actuarial profession? 8
3. What are the top AI risks for actuaries? 15
4. Reshaping the profession and the future actuary 20
5. The bigger picture: Increasingly capable AI, systems risks, and the future of the profession in a volatile world 25
References 26
Appendix A: Methods 27
Appendix B: Interview questions 29

Authors

Kun Zhao, MIT AI Risk Initiative, MIT FutureTech

Alexander K Saeri, MIT AI Risk Initiative, MIT FutureTech

Peter Slattery, MIT AI Risk Initiative, MIT FutureTech

Neil Thompson, MIT FutureTech

Funding and acknowledgments

This project was funded by a gift from the Society of Actuaries (SOA) Research Institute. The authors thank Dale Hall and Ryan Silhavy at the SOA Research Institute for their valuable input with project planning, recruitment, and review, and for their support throughout this project. The authors are grateful to all senior actuarial leaders who contributed their time and insights to the interviews.

Generative AI use

Large language models were used to assist with (1) report editing, (2) cross-checking themes and quotes against original transcripts, and (3) building a qualitative data analysis tool for organizing interview transcripts and for enabling thematic coding. All themes across interviews were coded manually. All content was reviewed and verified by the authors, who take full responsibility for the final report.

Suggested citation

Zhao, K., Saeri, A. K., Slattery, P., & Thompson, N. (2026). *AI Risks in the Actuarial Profession: Insights from Senior Leaders in Phase 1 of the Foresight Project*. MIT AI Risk Initiative, MIT FutureTech.

Contact

Kun Zhao (kunz1@mit.edu)

Executive summary

Overview and purpose

The Foresight Project is a multi-method and multi-phase research initiative led by the AI Risk Initiative at MIT FutureTech and the Society of Actuaries (SOA). It examines current and emerging risks to the actuarial profession, and its findings are intended to inform decisions about AI governance, professional education, and workforce development to prepare for the changes ahead. This report presents the findings from exploratory interviews with 13 senior actuarial leaders to understand the top AI risks for actuaries, and emerging trends and non-obvious themes according to senior leaders.

Major findings

AI use is experimental and fragmented in the profession: Senior leaders observed that AI adoption is highly uneven at the frontier, with widespread shallow use alongside pockets of advanced adoption. Actuaries are often caught between top-down pressure to adopt AI but no practical roadmap, restrictive early governance, and little bandwidth to experiment.

Top AI risks in the profession: Governance failure, transparency and interpretability, fraud and scams, privacy, and model failures including hallucinations were top concerns among actuarial leaders, converging with AI risk expert consensus. Minority views suggested that there is some shifting in risk priorities with maturing technology. Risks to professional reputation and human-centered risks of unemployment and loss of agency resonated more strongly with actuarial leaders than AI risk experts.

Reshaping the profession and the future actuary: There were deep and widely shared concerns around talent loss, displacement of actuaries, loss of fundamental skills for human oversight and accountability, and the relegation of actuaries to peripheral tasks as other fields move across professional lines. Actuarial leaders shared similar views on skills to preserve and promote: actuarial first principles, programming and coding understanding, as well as softer skills in communication and innovation. They also diverged regarding the deeper areas of expertise that could be owned by the profession: technological expertise, model-building, higher-level actuarial judgment and interpretation, and actuarial applications in complex systems.

The bigger picture: While AI's impact on the actuarial profession has been gradual, senior leaders anticipated disruptive changes ahead with increasingly agentic and autonomous AI, as well as broader societal shocks that could destabilize insurance markets. This was framed as a moment for the profession to draw on its reputation and lead the way in the safe and ethical use of AI.

Senior leaders' views on opportunities for professional bodies

- Establish structured guidance and standards for AI use and implementation in the profession
- Reform exams and education to strike a balance between foundational actuarial understanding (e.g., first principles) and reducing rote memorization
- Broaden the curriculum and consider different career tracks to accommodate a wider range of skills in systems thinking, higher-level actuarial judgment and interpretation, technological expertise, and model-building expertise
- Lead the way in championing the values that govern safe and ethical AI use in the profession, through education reform and public position statements

Directions for the project's next phases

- As AI use is varied and fragmented, identify different patterns and clusters of AI use rather than capturing averages or global trends across the profession
- Collect quantitative data and review the existing literature to test areas of consensus and disagreement among senior leaders, including differences by seniority, movement and displacement of actuaries, actuarial skills that are being eroded, and skills that are needed
- Monitor and track perceptions of risk using established taxonomies and measures

1. Introduction

Purpose of this research

The actuarial profession, like much of society, is entering a period of significant change, as advances in generative and agentic AI raise questions about how it will adapt. The Foresight Project is a multi-method and multi-phase research initiative led by the AI Risk Initiative at MIT FutureTech and the Society of Actuaries (SOA). It examines current and emerging risks to the actuarial profession and how AI is reshaping the profession. The findings provide an evidence base for decisions about AI governance and workforce development to prepare the profession for navigating changes ahead.

This report presents findings from the project's first phase, a set of exploratory interviews with senior actuarial leaders to ask "what are the top AI risks for actuaries", and to identify emerging trends and non-obvious themes. The focus on senior leadership aims to capture strategic insights for refining lines of inquiry and paving the way for a survey of SOA members in the project's next phase.

The report is structured in five main sections. The Introduction covers existing knowledge, relevant conceptual frameworks, and methods. The next chapter sets the scene for understanding the factors impacting AI adoption in the actuarial profession from a socio-ecological perspective. This is followed by two chapters on the top AI risks for actuaries from different angles: (1) risks *in* the actuarial profession and (2) risks *to* the actuarial profession and its professional identity. The concluding chapter situates these findings within broader societal systems and horizons, considering how mounting risks with increasingly capable frontier models could impact the profession's future.

Scope and caveats

What's in a name? Definition of AI in this report: Actuaries have applied a range of computational and statistical techniques now often grouped under the banner of AI. Throughout this report, references to AI denote generative and agentic AI and large language models (LLMs), unless stated otherwise. We focus on these in the report because of their general applications within actuarial work, labor market and societal impact, and the strategic uncertainty and risks that they pose. However, where relevant, narrower machine learning and deep learning tools are singled out given their significance for actuarial work.

Validation of participant perspectives: This exploratory study aimed to communicate insights directly from senior leaders, with minimal reference to wider empirical evidence. These views will be explored and tested in a research survey of SOA members.

What we already know about AI and risks in the actuarial profession

This project builds on existing SOA research on the actuarial profession, connecting this with MIT research on AI risks and its labor market impacts. Together, these sources capture where the profession currently stands and offer ways of framing how its future might play out.

SOA research on AI in the actuarial profession: The SOA Member AI Survey (Alaimo, 2025) found widespread AI awareness among more than 500 actuaries but shallow practical integration. The majority of respondents spent under 5 hours a week using AI/machine learning tools, with usage dominated by peripheral tasks. Engagement also varied with experience, with senior actuaries more willing to experiment with and recommend AI tools than their early-career counterparts. A subsequent expert panel discussion (Poon-Affat, 2026) examined how the profession thinks it should adopt AI and the most important traits to cultivate. This discussion emphasized professional accountability and cautious adoption, focusing on human-in-the-loop governance and early-career opportunities to develop human judgment and critical thinking.

MIT research and frameworks relevant to interviews: Recent MIT research can complement SOA knowledge and guide thinking on where the actuarial profession is headed. To understand how actuaries perceive AI risks, we drew on the Domain Taxonomy of AI Risks from the MIT AI Risk Repository (Slattery et al., 2026). This organizing framework classifies harms from AI into 7 domains and 24 subdomains, such as discrimination and toxicity, privacy and security, and misinformation (see Figure 1). These have been prioritized across various sectors by 272 AI risk experts through a Delphi consensus method (Saeri et al., 2026). Finance & Insurance was rated as one of the most vulnerable sectors, especially to fraud and scams, and discrimination, due to the potential for AI to be deeply embedded in critical decision making and where failures have immediate large-scale consequences. These findings were shared with interview participants to elicit views on AI risks in the profession.

Finally, the interview design in this project was guided by MIT research on how automation reshapes a profession (Autor & Thompson, 2025). The labor market impact of AI adoption depends on the extent to which automation removes routine, lower-expertise tasks (thus raising the bar for remaining workers) or removes more expert tasks (thus lowering the barrier to entry from adjacent professions). This introduces questions of what remains when AI absorbs some of the technical and analytical work making up the actuarial profession and the opportunities to take on new forms of expertise. This research was discussed in some interviews, to capture whether similar scenarios were unfolding in the actuarial profession.

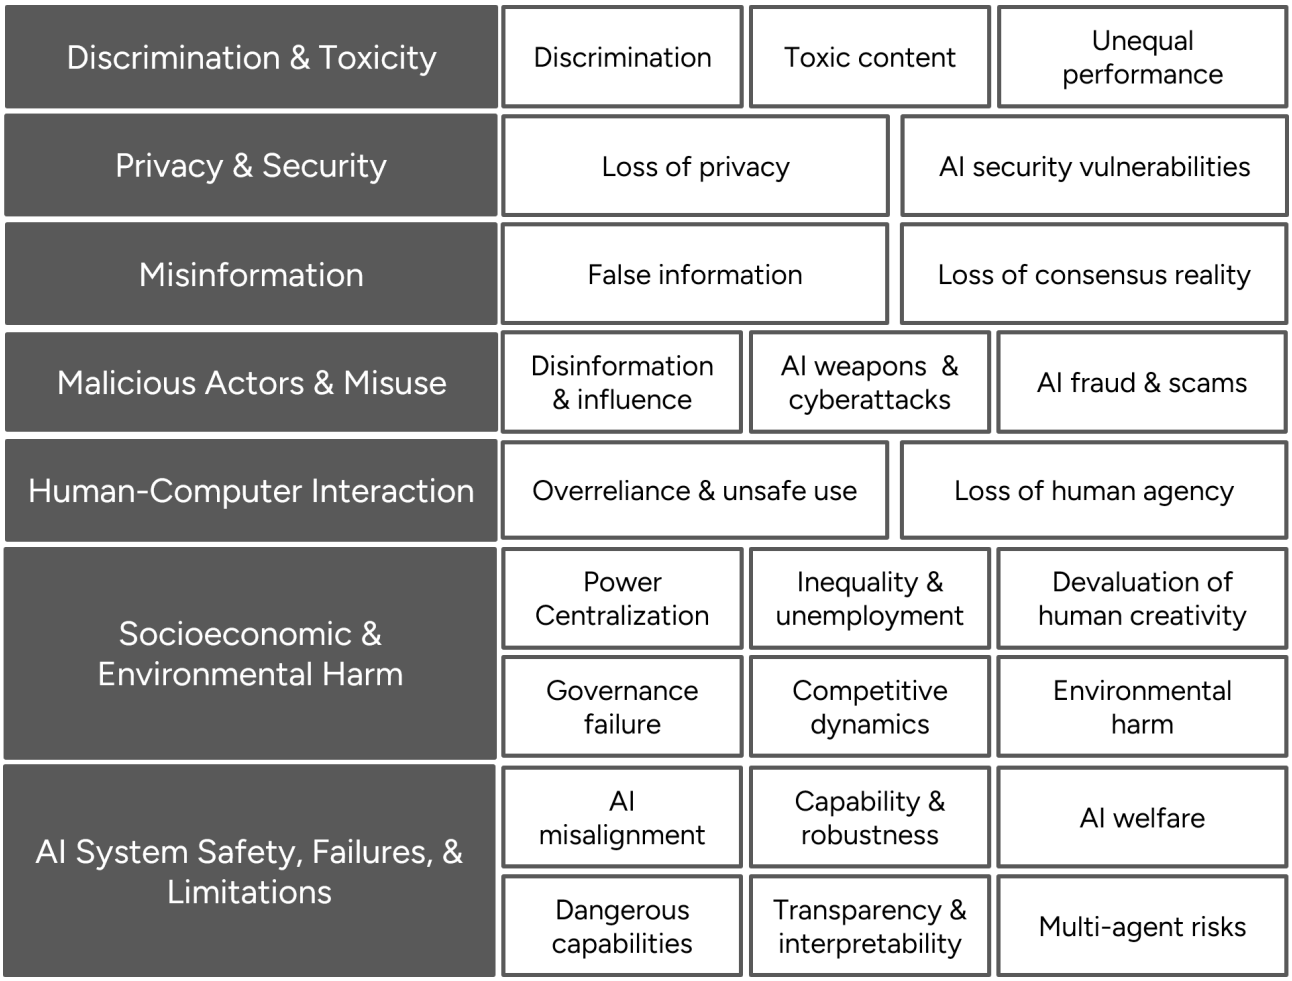


Figure 1: The Domain Taxonomy of AI Risks from the MIT AI Risk Repository (Slattery et al., 2026). *Grey boxes indicate domains of AI risk; white boxes indicate subdomains.*

How we identified and interviewed senior actuaries

The interview method and interview guide are provided in Appendices A and B. Participants were 13 senior actuarial leaders representing insurance carriers, Big Four professional services firms, actuarial technology firms, actuarial consultancies, global reinsurers, academia, and professional bodies. Geographically, about half of participants were based in the United States, with the remainder in Canada, India, South Africa, Australia, and the Asia-Pacific. One-hour semi-structured interviews were conducted online by the MIT research team, which consisted of questions about trends in AI use in the actuarial profession, the types of risks participants were concerned about, reflections on existing SOA findings and interview themes, and emerging insights.

2. What influences AI adoption in the actuarial profession?

AI use is both varied and deepening: from non-users to pioneers, and from everyday productivity to agentic workflows

There is wide variation in AI use in the actuarial profession and this was mirrored among senior leaders themselves, who ranged from having no experience with AI to being founders of AI-assisted actuarial modeling platforms and the field's leaders in AI strategy.

Leaders reported a range of AI tools, including general-purpose LLMs (e.g., ChatGPT, Claude, Gemini, DeepSeek), coding assistants (e.g., GitHub Copilot, Claude Code), AI-powered and agentic research tools (e.g., Perplexity, OpenClaw), custom-built agentic systems, local open-weight models, and internal/enterprise LLMs. While common day-to-day applications were writing, research, and administrative tasks, leaders reported increasing use of coding assistance (e.g., re-writing Excel spreadsheets in Python, converting legacy actuarial models, building actuarial tools and calculators), and along with it, a mainstreaming of technical skills. Retrieval-augmented generation (RAG) and agentic RAG pipelines were cited as valuable for regulatory knowledge access and compliance work. A recurring theme was the transition from AI as a productivity tool to agentic systems capable of autonomous multi-step workflows (e.g., executing end-to-end legacy model conversion, retrieving data and building presentations).

Senior leaders report experimental and fragmented AI adoption

Senior actuarial leaders painted a picture of the overall profession in an exploratory, experimental, and fragmented phase of AI adoption, where approaches are neither formalized nor standardized. There was tension between leadership enthusiasm and on-the-ground caution, with one leader contrasting the fact that "AI has just become accepted without being proven" against the extensive testing required for traditional systems development.

> *"I don't see successful examples that we can say, aha, this is how someone has solved this big problem or changed workflow in some way or made some advance that then the whole profession can follow. What I see is just little experiments here and there." – Senior leader, on AI adoption*
>
> *"I'd say the starting point here is that companies are not at the point where there's been widespread sort of standardized adoption of AI. It's more, I'd say, and it's still in that kind of experimentation phase [...] I'd say is that it is very much still in the exploratory phase for most companies" – Senior leader, on AI adoption*

Despite recent developments in generative and agentic AI, adoption was uneven at the frontier, with leaders describing a pattern of widespread but shallow use that did not indicate a qualitative shift almost a year on from the SOA Member AI Survey:

> *"If I were to take a poll of [5000 consultants] and say, you know, how many of you have used, you know, GPT this week, it would be, like, 95%. But if you kind of step into, like, how many of you have incorporated some kind of AI tools into your process, I'd bring it down to 25%. How many are using Claude Code for writing PowerPoints? It's, like, me and three other people. So I think it kind of quickly falls off in terms of who's at the frontier." – Senior leader, on variation in AI use*

Leaders also cited positive case studies of insurers investing in AI adoption and governance by establishing dedicated AI leadership, funding capability building, standardizing use cases, experimenting, and testing models for bias.

Senior leaders describe cautious and mixed views toward AI adoption

Despite prevailing optimism and pressure to adopt, some leaders raised concerns about returns on investment, the limits of current tools, and the crowding out of other opportunities to innovate.

Doubts about the return on investment and the cost of agentic AI:

> *"But I think the biggest thing is people are going to start to find "I can't throw agents at everything because it's actually becoming way too expensive", right? The amount of money I'm burning on processing time and tokens [...] So I think there's a piece where companies are going to have to really start thinking about what's the right governance model of when I use an agent, who's got access to it. [...] And I just don't think companies have gotten to that point just yet to figure out what the cost benefit equation is going to be, by use case, technology, type of agent." – Senior leader, on what needs to be discussed more*
>
> *"Sometimes they say that some organizations have received ROI from investments, but I also see the other one, like 90% of the organizations are yet to see any significant returns from investment." – Senior leader, on uncertainty about whether investments in AI will pay off*

Productivity illusions and the "last mile" of quality: Leaders cautioned that a polished demo or minimum viable product is easy to produce but production-grade quality is harder to reach: "First you have an idea [...] and then you look into that, there's holes everywhere. And if you are trying to fix those holes, you may create new holes everywhere." Similarly, another leader reported superficial results when using generative AI tools to build a prototype for researching the mechanics of solvency: "If I wasn't an actuary, I would think that it's great. But being an actuary, I'm like, eh, it doesn't look great. [...] And that was just faster and more straightforward and a lot less frustrating [to read the regulation myself] compared to having to triple check everything that AI says."

Opportunities to innovate in narrow AI tools as well as developments in actuarial-aware predictive modeling: A separate concern is that enthusiasm for automating actuarial work with generative AI is crowding out innovation in predictive modeling (e.g., paradigms such as ICEnet, Individual Conditional Expectation network, which injects actuarial requirements in the design and parametrization of deep neural networks for making predictions on actuarial tasks). Leaders also pointed to new developments in narrow predictive tools beyond LLMs, including deeper machine learning methodologies and a new frontier of transformer-based tabular foundation models (e.g., TabPFN) that have clear potential to benefit the profession.

> *"Right now, everyone's very interested in generative AI. How do you use LLMs? Can you automate actuarial work? And I think that's all valuable, but it shouldn't be at the cost of surrendering one of the core disciplines of actuarial science, which is building good predictive models." – Senior leader, on the biggest gaps*

AI adoption is driven by individual, professional, organizational, and societal factors

Senior leaders described many influences shaping AI adoption, and it became clear early on that the professional role was only part of the story. Individual readiness, organizational pressure, industry regulation, and pace of change all interacted to influence how leaders viewed the opportunities and risks of AI, and how they saw the profession adapting to it. For this reason, a socio-ecological framework was used to organize and distinguish between nested layers of influence, i.e., the individual actuary, the actuarial profession, the organization, and society (Bronfenbrenner, 1979; Nguyen-Trung et al., 2025). These influencing factors are presented in Table 1, which also summarizes leaders' views on their impact on AI adoption, along with illustrative quotes. Influencing factors that leaders agreed on are denoted by an upwards (driver of adoption) or downwards (barrier to adoption) arrow, and those that had conflicting views among leaders are denoted by a yellow circle.

Table 1: Major influences on AI adoption identified by senior actuarial leaders

Influencing factor	Influence on AI adoption	Quotes from senior actuarial leaders
Individual-level factors		
Experience and seniority 🟠	Conflicting views on whether seniority was associated with AI adoption, with some noting that junior staff with "boots on the ground" were more cautious while, according to others, senior staff were using AI more critically	"There's basically a lot more optimism about it on a senior level compared to a junior level. I attribute it mostly to the fact that the junior level has probably already experimented with it and kind of saw the same issues that we've seen and said, look, you know, that's one way to lose my credentials." "The senior folks? Actually they're using it, I believe the right way that I'm proposing to use it. They're getting their hands dirty [...] they've done the work [...] [And they're seeing] the difference in the output in terms of what they've worked on and what the AI produced."
Personal appetite for technology 🔺	Individual appetite and existing experience with technology were stronger drivers of AI adoption than factors such as practice area, type of work, or seniority. Two actuaries at a similar career stage could have vastly different experiences with AI adoption.	"There was a set of people who are very excited about AI and who's using these tools every day. And there was another set of people who were completely blindsided by what we were talking about. And those are very smart actuaries, right? They are doing very well in the actuarial profession. So, I do feel that, you know, there is a personal preference to people whether they want to get involved with tech or not."
Profession-level factors		
Adaptive capacity 🟠	Disagreement around whether lack of adaptability among actuaries was a barrier in itself. A recurring theme was the fact that actuaries have already evolved through several waves of technological change, though this may reflect survivorship bias.	"Actuaries have been pretty good at staying on top of the latest developments to help us do our job, you know. [...] And, you know, I think we're open to that sort of adaptability. I think we do pretty good at that. But this is a whole new level. So, we're kind of on edge." "I mean, we see a lot of progressive actuarial functions. We see, you know, some actuaries that really want to be super progressive on things." "I know there were actuaries who didn't survive that transformation and who crashed and burned and left the profession and didn't make the change. And I think with AI, we're probably in the same situation today that we were in in 1984 when Lotus came out - we're going to have to transform what we do."

Influencing factor	Influence on AI adoption	Quotes from senior actuarial leaders
Risk aversion	The risk-averse and procedure-driven nature of the profession is a barrier to AI adoption and a risk to the profession's survival	"We get so afraid to try something that we try nothing, and then the profession gets stagnant. And it can be, I think, potentially it's a product of the way that we train actuaries. We train them to regurgitate and report a list of things, to get to an answer. Here's the number. The exam doesn't go to 'Okay, so you got this number, what does that mean? What will you do with it?'" "I think a lot of people are attracted to the actuarial profession because they sort of view the process as this meritocracy. They put in their five to ten years, they pass these exams, and they're delivered a relatively well-paying job [...] But a lot of actuaries are probably not [more innovative]. What keeps me up at night is I think all those people are in for a rude awakening in three to five years."
Professional standards and accountability	Strong code of ethics, professional standards, professional accountability and requirements for human sign-off are simultaneously barriers to rapid AI adoption as well as a means of managing AI risks	"It doesn't matter what system you use to produce the results, the licensee is actually responsible for all the outcome or the results that are produced. And so therefore people get very nervous when you introduce any AI or non-deterministic components in the process." "There's no other profession that has a process as rigorous as that of the actuarial profession. So, with that, ask the public, who would you trust more, you know?"
Tacit knowledge	AI can assist with knowledge that is explicitly documented but much of valuable actuarial expertise exists as undocumented tacit knowledge accumulated over years of practice	"Because we are full of, you know, 30 plus years of context. And it's really hard to, like, kind of jam that context. I'm not able to upload my own vector database into this thing"
Organization-level factors		
Internal pressure to reduce cost ▲	Perceived cost savings and efficiency were strong drivers of AI adoption. One leader observed that the dominant motive for adoption was "cost arbitrage" but questioned whether headline savings would materialize, noting that innovative companies were instead shifting toward customer-focused, growth-oriented uses of AI	"There is a huge expectation for high productivity and cost cutting [and] you are putting so much pressure on your workforce to do this."

Influencing factor	Influence on AI adoption	Quotes from senior actuarial leaders
External market pressure ▲	Revenue and client and investor pressures were strong drivers of AI adoption, where simply signaling AI capability has market consequences for insurance companies	"It's mostly market-driven because businesses themselves aim to stay competitive. [...] They want to attract as many profitable customers, for instance, as they can. And so, you know, the business really drives that. And I think actuaries have been fairly good at responding to them, but we've always sort of been the users of the tools."
Security and IT infrastructure ▼	Organizational and IT security policy and infrastructure were strong and at times self-defeating barriers to adoption, pushing staff to riskier or more expensive workarounds	"Oftentimes, it's the corporate [that] holds them behind, it's IT that holds them behind, it's security that holds them behind, and sometimes it's the regulations, quite frankly, that hold us behind. [...] [They] don't have access to Gemini, as an example, right? So, they're doing stuff on their own laptops, and obviously that introduces data security risks. They can't move the datasets outside of the company, so they probably can't do the analysis."
Actuarial teams de-prioritized	Under pressure to demonstrate efficiency or returns, organizations direct AI investments toward other teams before actuaries	"Now, because of the pressure every CEO feels to invest in AI and show results, if they have a few million dollars to invest, they are not going to the actuarial team, they're going to the tech team. And the tech team kind of uses actuaries as consultants to just ask a few questions and they design their own stuff."
Lack of bandwidth ▼	Lack of resources and capacity for investing in AI adoption due to business as usual	"Everyone's under a lot of workload [...] So therefore do you even have the bandwidth from the subject matter experts to actually innovate in that area or not [...] If you are already 90% loaded, that 10% doesn't help you to clear your mind and be creative."
Lack of implementation guidance to adopt AI ▼	Absence of firm guidance, examples, or a standardized pathway, with adoption confined to isolated pockets rather than a coordinated approach	"Possibly my CEO is saying, use more AI in your work, but, you know, she or he doesn't have any kind of firm advice for how to do it." "There's no widespread standardized implementation of these tools"
Organizational size	Smaller organizations able to develop informal workflows faster and operate with greater trust	"And so for [smaller teams], it's a little bit easier to trust each other and know each other personally, to know that people would not do shady things. So, in bigger organizations, it is certainly a bigger problem because you have new employees, you have pressures from management."

Influencing factor	Influence on AI adoption	Quotes from senior actuarial leaders
Organizational appetite for risk		"All the big companies don't all move at the same speed. And so that's maybe more company culture. Some companies are more willing, it comes really, what it comes down to is risk appetite"
Society-level factors		
Insurance industry moves slow ▼	Leaders characterized the insurance industry as a slow and cautious adopter, partly owing to heavy regulation	"And probably five, ten years ago, banks were looking for the types of roles that insurance companies are now looking for. So, I think insurance companies are generally behind the curve." "I mean, the insurance industry by itself is not a fast adopter of anything, right? I mean, if anything, we've got a bad reputation for, you know, being a sleepy industry."
AI moves fast ▲		"How do we as a profession, let alone an industry, let alone all the other companies, keep up with the pace of change, right? The decision I take today is redundant tomorrow, and I still need to make progress in meeting my quarterly returns and, you know, seeing the investment I'm putting into modernization, right?"
Regional differences in regulatory requirements and AI adoption trends ●	Leaders noted regional variation in the pace of AI adoption and regulation, but fewer qualitative differences emerged across geographic lines	"Because of the release of [DeepSeek previously], the Chinese insurers are moving fast [...] They have the hardware and build DeepSeek – their localized AI system, but it looks like at the current stage, North America is catching up."

▲ Agreement among leaders that this was a clear driver of AI adoption ▼ Agreement among leaders that this was a clear barrier to AI adoption ● Conflicting perspectives across leaders

3. What are the top AI risks for actuaries?

Senior actuarial leaders were asked about risks from AI in two ways:

- What risks from AI they were most concerned about
- Their reflections on the top AI risks for Finance & Insurance as rated by AI experts in previous MIT research (Saeri et al., 2026; see Figure 2)

Leaders' responses overlapped with risk expert consensus in some places but diverged in others and there was some indication that risk priorities were beginning to shift as the technology was maturing. Overall, leaders shared the top concerns with AI risk experts for risks such as fraud and scams, governance failure, privacy, and transparency and interpretability. Leaders were relatively less concerned with AI weapons and cyberattacks with respect to the profession. Common themes among actuarial leaders, but absent or lower-ranking in AI risk expert consensus were concerns with reputational risk and human-centered concerns involving overreliance, loss of human agency, and unemployment.

Traditional governance and risk management approaches failing to address AI

Governance failure, where AI capabilities and adoption outpace the capacity of organizations to manage these risks, was among the top priorities judged by AI risk experts. Among actuarial leaders too, risk of governance failure came up in almost every interview. Senior leaders reported that traditional actuarial governance and risk frameworks are built for routine, repeatable, and predictable output. They are not merely lagging but are inadequate for generative AI and LLMs, where output is non-deterministic and models are continually updated and deployed against moving targets.

> *"Governance and risk management is the one where I feel a bit nervous. Because if things are mechanical and every quarter you produce the same thing again and again, then I guess the governance and risk management frameworks that are in place work. [...] I get very nervous when we get into the non-recurring or the non-regular work, whereby everyone is testing for the first time, you can call it PoC you can call it pilot, you can roll out a limited... when is it really in production? [...] So when do you take the result for granted? And when is it that things are locked down and in production?" – Senior leader, on governance limitations*

This is further complicated by the rapid pace of change. As one leader put it, by the time a framework is written down "the world's going to have changed", rendering it near-irrelevant on arrival. Leaders raised further issues with model risk management: as models grow more specialized, the small pool able to validate them ceases to be independent. As agentic AI is increasingly used "like a person", even defining a model and maintaining a model inventory becomes difficult.

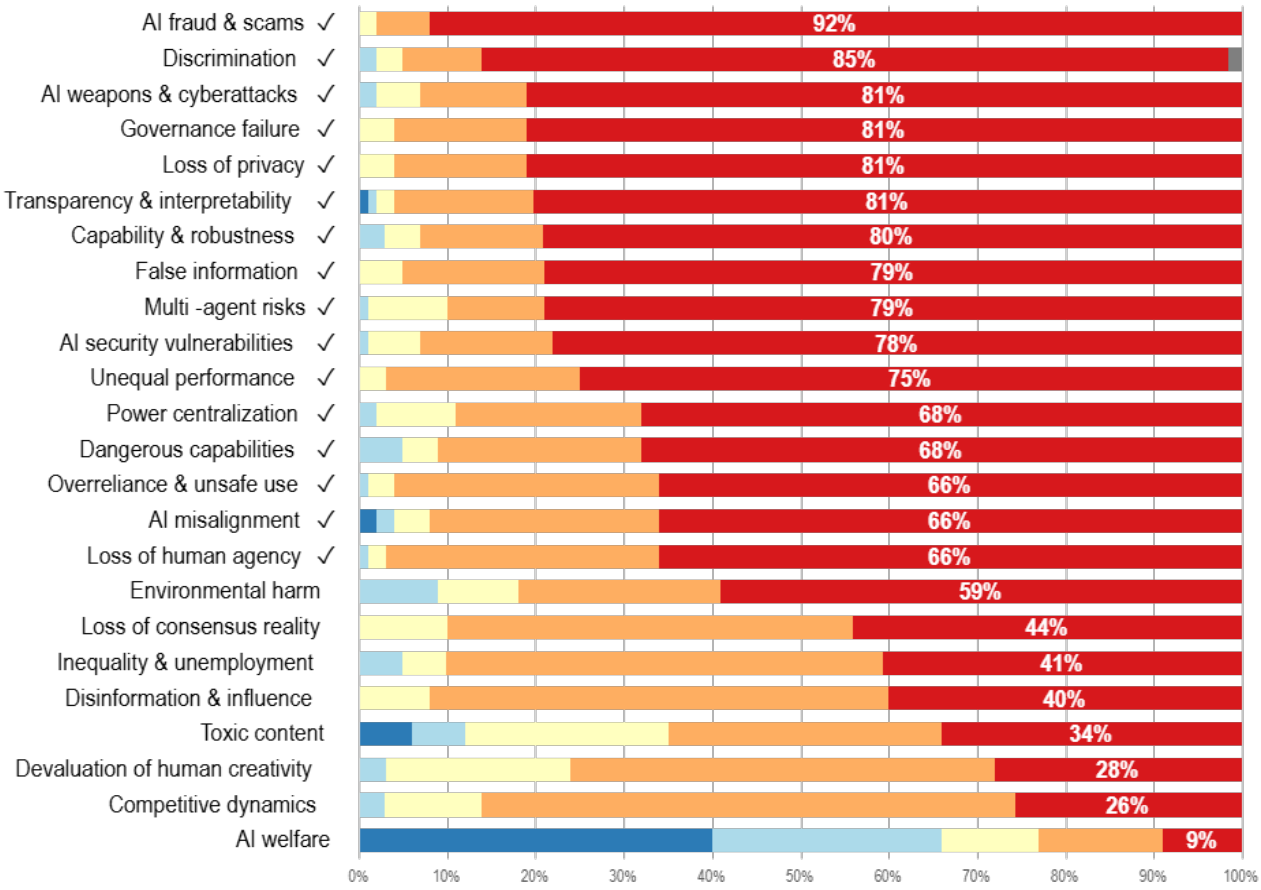


Figure 2: AI risk experts' ratings of vulnerability of the Finance & Insurance sector across 24 AI risks (Saeri et al., 2026). *Checkmarks denote statistical consensus in a sample of 272 AI risk experts. Colored bars indicate percent selecting "Not at all vulnerable" (dark blue), "Minimally vulnerable" (light blue), "Moderately vulnerable" (yellow), "Highly vulnerable" (orange), and "Extremely vulnerable" (red).*

Inability to interpret model outputs and audit model internals

Almost all senior leaders highlighted concerns with transparency and interpretability, which resonated strongly with the profession and were also highly prioritized in AI risk expert consensus. Concerns clustered around the inability to reproduce or validate AI outputs under existing model risk management frameworks and the shift from deterministic to probabilistic models, which is fundamentally incompatible with requirements for actuarial sign-off and professional accountability.

> *"Our professional standards says that our results should be replicable. And they should be explainable, right? And if I don't do that, I'm at risk of losing my license because, [...] I need accountability, right?" – Senior leader, on risks of errors*

AI-enabled fraud and scams, and discrimination

Together, fraud and scams and discrimination represented the subdomains where the Finance & Insurance sector was judged by AI risk experts to be most vulnerable. Fraud and scams were acknowledged as prevalent concerns among senior leaders. In contrast, there

were mixed views on risks associated with unfair discrimination. Some leaders believed that it was overstated as AI merely mirrors pre-existing biases ("It's already there [...] you're just repeating what's happening, right?" and that "insurance as a business is discriminatory").

Loss of privacy, AI system security, and hallucinations or other model failures

While most leaders flagged privacy and security as dominant risks, a minority view was that they are less prominent than in previous years. One reflected that previously the concern was over what could securely go into the AI model but these days they were more concerned about overreliance on AI. Others reported that privacy and security concerns are starting to be addressed with new ways to contain and govern AI tools within secure organizational environments.

Likewise, while most leaders viewed hallucinations as a primary risk, others reported that they are a declining concern and "not the sort of problem that it was a couple of years ago", due to developments in the safeguards and workflows around tools.

However, another leader flagged that multi-agent risks should be ranked higher, describing cascading or compounding errors that come about from interactions across networked agents, and the expectation that they will increase.

Reputational risks: Loss of trust in the actuarial profession due to underestimation of AI-enabled harms

A common theme among actuarial leaders, but absent from AI risk expert consensus, was the risk to professional reputation. Leaders worried that hard-earned trust in the profession could be lost when insurers are unable to meet their obligations, whether due to AI misuse, failure, or wider macroeconomic disruptions:

> *"So, the really big risk for us is these organizations within, you know, within our economy failing to meet their obligations because we're in charge of making sure that we have the right finances set aside to meet those obligations. So that's like our key pillar, that's what actuaries are valuable for if nothing else. So, the risks would be that relationship between, we the profession and the people over the country, the world, the risk is that that relationship is harmed, and it can be harmed in a number of ways." – Senior leader, on the most concerning risks*

Human-centered risks: Lost agency, overreliance, and labor displacement

Senior leaders were also drawn to risks concerning the human cost of AI adoption, which were lower-ranking in risk expert consensus. These included loss of human agency and autonomy: "But it's this dehumanization, I guess, of humans in the great shape of things and

loss of agency, loss of control, loss of judgment, humans becoming more, in a way, less creative". There were also deep concerns about the risk of unemployment and displacement of actuaries, explored in detail in the next chapter.

> *"And then unemployment and equality, I would say should be much higher [...] you can see a lot of the big tech firms are chopping people, letting people go and so forth. [...] At least on the life side, I understand that the actuaries are ... finding job opportunities a bit hard to come by." – Senior leader, on prioritization of risks from AI*

Emerging AI governance and risk management approaches among actuarial leaders

Leaders described drawing on several overlapping layers of governance, ranging from contractual controls and organizational frameworks to technical standards, regulation, and, more generally, professional accountability through codes of conduct and standards.

Contractual and organizational controls: Guardrails written into contractual relationships, internal organizational frameworks (e.g., "Do's and Don'ts of AI" on a company website)

Technical and assurance standards: Established frameworks (e.g., SOC 2, now straining under generative AI), NIST Generative AI Profile, ISO 42001

Regulation: NAIC initiatives, state insurance regulations, OSFI's model risk management framework (Guideline E-23)

Professional standards and codes of conduct: The profession's disciplinary body (Actuarial Board for Counseling and Discipline), professional accountability and sign-off, "timeless" actuarial precepts and codes of ethics that "applied before the computer, and they continue to apply now [which] hold up pretty well, even under AI". A view shared by almost all was that professional accountability due to timeless actuarial precepts, codes of ethics, and professional standards was a general protective factor:

> *"And I think trust at the end of the day is key. AI will not sign a report, will not be sued, will not be asked to, you know, to be disciplined by a judiciary body, like an actuarial body, an actuarial discipline process. I think what protects the actuaries' profession is, all of these things, the examination process, the code of conduct, the standards that we have to rely on, and our disciplinary process, which puts the onus on the actuaries to do good work." – Senior leader, on factors that are protective of the profession*

Informal human-in-the-loop and validation practices:

- Task decomposition and breaking down AI work into smaller verifiable chunks
- Verifying using deterministic portions of the model (where AI is readily tested but adds little value) and moving from here to more subjective portions of the model (where AI adds more value but verification is less reliable)

- Model consensus or ensemble paradigms, where many models are used and the average is taken to examine whether it can reproduce the results of one model
- Using non-deterministic AI tools to build deterministic workflows, e.g., AI coding agents to redesign data pipelines with human oversight, leaving behind fully automated, deterministic workflows that run without further AI involvement

Senior leaders cautioned that actuaries needed to be actively vigilant and take on a preventive mindset when working with AI: “I always tell people, treat it like an intern and interns tend to make a lot of mistakes”. However, the unpredictable and counterintuitive nature of LLM errors posed further challenges, with others noting that they were fundamentally different from the mistakes of junior actuaries (e.g., needless model complexity, surface-level fixes, resurrecting old errors after manual fixes), forcing new approaches to error checking.

4. Reshaping the profession and the future actuary

This chapter focuses on a collection of risks and proposed actions concerning the actuarial profession and professional identity itself: the talent pipeline of junior actuaries and preservation of core skills, displacement of actuaries across all levels, and the redrawing of boundaries with adjacent professions.

AI has multiple impacts on the junior talent pipeline: talent loss, displacement, and diminished learning opportunities

There was near-unanimous concern over the pipeline of junior actuaries. This centered on two distinct but compounding challenges: the loss of entry-level work and the loss of hands-on opportunities to build fundamental understanding and judgment, the very skills needed for human oversight and accountability. As one leader put it, the worry is not only what happens to a team replaced by AI, but how anyone plans for "career advancement where the machine is doing all the core learning".

Further, one leader described the exploitative nature in which human expertise is "downloaded" into automated workflows but staff behind them are not retained:

> *"People are very skeptical now - it's like my boss has asked me to write down what I've done, what are my thinking processes [...] I'm hearing this slightly more, and then they will think 'Am I being replaced by AI eventually because they just want to download my brain?' [...] It's actually quite depressing to that person. So how do you manage the psychological aspect?" – Senior leader, on non-obvious themes*

Leaders also observed talent loss to other fields as the actuarial curriculum struggles to keep pace:

> *"[Y]ou take these juniors and know that they're very interested in AI, very interested in code, very interested in machine learning, and you actually equip them to build their career in that direction from the get-go. So, if you're asking new junior actuaries to do what junior actuaries did 10 years ago, you might find that they're not really interested in doing that anymore. And the profession becomes a little bit less interesting as a career choice [...] Last year at the biggest university here [...] where actuarial science is extremely prestigious, [it was] the first time ever it's been edged out by computer science." – Senior leader, on the talent pipelines of junior actuaries*

At the organizational level too, leaders observed movement to other areas as career opportunities and investment favored AI-focused roles over traditional actuarial functions:

> *"[My engineers] have realized that if they remain in my team, their career path is very different because it's an actuarial team, right? So, they are actually going to AI teams that are focusing on actuarial. I am losing my best people to tech teams now because they think that they will*

> *get a better career if they remain within the AI team and then support them with actuarial" – Senior leader, on issues that need to be discussed more*

AI is intensifying longer-running encroachment on actuarial work by adjacent professions

Other leaders observed that AI adoption is redrawing professional boundaries and is accelerating a trend predating generative AI, with data science, statistics, and finance moving into actuarial domains. They presented multiple examples in which "tech [is] cannibalizing risk areas" and building systems themselves, which erodes opportunities for risk professionals "unless they immediately start getting better at technology and AI from a code point of view". Another leader similarly noted that automation is enabling "statisticians, claims guys, other people [to come in] and [do what] we're doing because now they've got AI agents that can do a lot of what the actuaries used to do".

Other leaders predicted that these are contributing to the disappearance of certain insurance markets (e.g., personal lines insurance) and the "downgrading" of actuarial tasks as the field becomes dominated by data scientists and other professions: "[F]rom a professional perspective, it's quite sad that actuaries are almost being relegated in some markets to do different tasks, or tasks that used to be done by actuaries and are no longer being done by them."

AI's impact varies across actuarial roles, devaluing some expertise while deepening other specializations

When reflecting on labor-market scenarios based on MIT research concerning the redistribution of tasks and shaping of expertise, one senior leader reminded that:

> *"[B]eing an actuary is not a single homogenous activity. There are many different types of actuaries and within each one of those types, I do think there'll be a different impact [...] If we go to the one extreme personal lines pricing in the UK, I think that change has already happened or is already happening. And that's very firmly the taxi driver bucket [i.e., expertise devalued by technology]. We've got new people who can apply new skills like getting a GPS machine or wielding tools from data science and actuaries have become displaced. On the other hand, very much in the life insurance space [...] I think actuaries are responsible for product pricing and for product valuation. Those sorts of people are still very much doing the jobs they always did. It becomes a little bit more niche [i.e., specialization is preserved]." – Senior leader, reflecting on research on automation and expertise*

Considering the scenarios above, leaders were optimistic that advances in AI could "liberate actuaries from drudgery" and move junior actuaries into more advanced work faster. Others agreed that "actuaries are growing up faster" and that enhancement through AI tools will see an acceleration of the actuarial learning curve with skills that previously took 5-7 years being developed in 2-3 years.

> *"Therefore, the actuarial skill set is going to be accelerated from the standpoint of early in your career, you're not going to distinguish yourself because you're able to build models better than somebody else or gather data. If that's all going to be normalized, where you're going to distinguish yourself is your ability to understand regulations faster, understand some of the insights, being able to communicate those more effectively. Those skills that usually come later in your career are going to be accelerated in actuarial development. And therefore, early-in-career actuaries are going to have to build those skills probably in a shorter time frame than they've had to in the past." – Senior leader, on changes in the talent pipeline of early career actuaries*

Intuitive understanding of actuarial first principles is most valued but hard to preserve with reliance on AI

Senior leaders agreed that intimate knowledge of actuarial first principles and foundational concepts is essential, but most at risk of being underdeveloped as AI absorbs the tasks through which that understanding and intuition are built.

> *"If you've done the math and done the spreadsheets enough times, you start to understand the drivers of things. You understand the sensitivities and the drivers and you start getting senses for OK, this is too high or too low [...] I think people have lost that a little bit over time as the models have done more and more of that work." – Senior leader, on the loss of core skills*

A similar bind applies to programming and, with it, coding literacy. Many agreed that the ability to understand and evaluate code is important, especially given the risk of overreliance on closed, "black box" systems: "You still need to understand all the logic, like the calculation logic, coding logic, coding structure [...] Otherwise it will be totally misleading and you don't know what you are creating. You trust AI too much, then that's a problem". This is distinct from coding ability, which is being commoditized: "Somebody's very good at coding - does it matter now? Probably less relevant because I can just get the AI to code for me." In both cases, there was no clear resolution on how comprehension and judgment could be maintained with automation of tasks through which foundational knowledge is developed.

Importance of soft skills for effective communication, as well as innovation and entrepreneurialism

Senior leaders saw value in capabilities that could not be easily replaced by AI: "soft skills, how to interact with other teams [...] communication, relationship skills", along with collaboration and negotiation, which "are going to set someone apart from having purely hard skills". So, too were inquisitiveness, innovation, as well as entrepreneurialism, which leaders acknowledged were weaknesses in the profession: "Actuaries too often just do routine things and get paid well and have very comfortable, secure jobs, which is why it's hard for them to strike out and do something new and different".

To make way for these, leaders suggested that, beyond mathematical fundamentals, some rote skills could be cleared out from the existing syllabus: “And if the system of the exam is so focused on regurgitation and memorizing lists, then you aren't building people to think and draw conclusions and say, ‘So what’?”.

Differing perspectives on what actuaries should be leaning into for deeper expertise – builders of models and bearers of judgment

Some leaders sought to protect actuarial expertise by building deeper modeling expertise and technological skills, and others prioritized the judgment and broader applications that AI falls short on. In the first group, some leaders emphasized that:

> *“And as a profession, we've always been the builders of our own models, and that discipline is rapidly becoming ignored. So, I think the biggest opportunity is to build good models, whether they're general or narrow, to take advantage of the latest advances in AI, in deep learning, also in the field of machine learning and reinforcement learning, take all of this incredible wave of advances and put an actuarial flavor on it. […] Actuaries should be innovating in this space. And that isn't really happening outside of a few very select groups of researchers globally. So, I think in a way we're losing as a profession, one of the most important things which we did was being the people who own and design and parametrize and operate our own flavor of models. That, to me, is without a doubt the biggest gap" – Senior leader, on the biggest issues in the profession*

Others stressed the importance of technological expertise and the role of professional bodies to develop this within the actuarial profession. One leader proposed a dual career track within the SOA to extend the credibility of the SOA “stamp” to technological skills. This would consist of a traditional path toward chief or appointed actuary and a parallel "actuarial technologist" route which combines foundational actuarial training with external technology certification and SOA-supervised projects.

Other leaders looked away from specialization in technology, programming, and model-building altogether, arguing that actuaries should lean into actuarial judgment in the form of “individual, experience-based, context-based” expertise, which AI cannot substitute for. Future actuaries should report not just what the numbers show but also derive the “so what”, and translate it into practical recommendations and foresight:

> *“The actuary is the one looking through the rearview mirror telling somebody how to drive the car […] But then you also kind of want to marry that with the actuary who's looking forward - seeing what is emerging. That could be medical advancements, technological advancements, demographic shifts […] kind of the combination of looking out the rear and looking at a crystal ball of what could happen.” – Senior leader, on actuarial judgment*

One leader argued that modern risks from climate change and AI carry complex societal and catastrophic consequences, yet are characterized by deep uncertainty and little historical data to model from. To deal with these, actuarial education should move beyond traditional

statistical skills to better equip actuaries with knowledge of geopolitics, economics, and systems change.

> *"We tend to take these people [with a high] score in math, we then immerse them in three years of data science and statistics and math, and then we turn them loose and expect them to be able to deal with climate change or AI or geopolitical risk [...] I would argue that you're doing a poor job preparing people, because they're not thinking about economics, about politics, they're just analyzing data. And data breaks down with AI, geopolitics, and climate, right? That's where your data gets you in trouble because the past is not predictive of the future" – Senior leader, on the importance of systems change*

Expanding the aperture of actuarial work

While some leaders felt that actuaries need to "consolidate and defend our territory a little bit more", others pointed to expanding domains that actuarial risk reasoning could be applied to: banking, energy markets, climate, political risk, and medical economics. Although current narratives have been dominated by automation and productivity use cases, one leader pointed out "tremendous opportunities" to explore the expanded risk landscape for actuaries to do "better financial risk analysis", "better forecasting analytics", "better risk scans" and "business intelligence".

5. The bigger picture: Increasingly capable AI, systems risks, and the future of the profession in a volatile world

This final chapter steps back to comment on broader societal impacts on the actuarial profession alongside developments toward artificial general intelligence. While this report opened with a picture of AI adoption in the profession as incremental and slow, most leaders foresaw greater change, with some remarking on seismic societal shifts.

> *"So, you know, currently it's not, it's not making so many headwinds that all of us are forced to say we must change everything that we do. Probably that moment's going to come, like, piece by piece. I bet there's going to be some really big bricks that fall." – Senior leader, on what is not being discussed*

One leader questioned the view that "[a] fun thing to say in 2023 was: you might not lose your job to AI, but you're going to lose your job to somebody that's using AI". They doubted that this "paradigm kind of holds anymore" as AI begins to generate "new AI capabilities on its own" in faster cycles, leading to the breakdown of human-in-the-loop conditions and "AI will be using AI tools".

A further theme was AI as a source of societal volatility. While most discussions centered on risks endogenous to the insurance industry, some leaders pointed outward to wider shocks (e.g., autonomous vehicles, pandemics, geopolitical conflict, and macroeconomic disruption) that could play havoc with assets and liabilities, undermine planning, and destabilize insurance markets. This is in line with MIT findings from a Delphi study of 272 AI risk experts, who estimated most AI risks as having at least 10% probability of catastrophic outcomes (e.g., more than 1 million deaths or more than USD 100B in financial losses) over the next five years (Saeri et al., 2026). The scale of what might be at stake drew one leader to reflect on parallels with the climate risk concept of "planetary insolvency", describing tipping points and catastrophic harms in scenario planning that push the planet or humanity past a point of no return.

Finally, one senior leader placed the profession at "the crossroad in our history and humanity where we could go, you know, the wrong way or the right way" and navigation requires clear moral direction. Given the actuarial profession's long-standing reputation and moral credibility, they argued a case for actuaries to become advocates for the ethical use of AI. This could take the form of exam and education reform, alongside declarations or position statements setting out the profession's stance and values on safe and ethical AI.

References

Alaimo, J. (2025, November). *SOA member AI survey: Summer 2025*. Society of Actuaries Research Institute. https://www.soa.org/globalassets/assets/files/resources/research-report/2025/2025-11-ait-172-ai-quarter-survey-report.pdf

Autor, D., & Thompson, N. (2025). *Expertise* (Working Paper No. 33941). National Bureau of Economic Research. https://doi.org/10.3386/w33941

Bronfenbrenner, U. (1979). *The ecology of human development: Experiments by nature and design*. Harvard University Press. https://doi.org/10.2307/j.ctv26071r6

Nguyen-Trung, K., Saeri, A. K., Zhao, K., Boulet, M., & Kaufman, S. (2025). SeCOM-B: An integrated model for understanding human behaviour change in wicked socio-ecological problems. *Socio-Ecological Practice Research*, *7*(3), 367-386. https://doi.org/10.1007/s42532-025-00227-y

Poon-Affat, R. L. (2026, January). *Actuarial mindsets for leading in the AI era: An expert panel discussion*. Society of Actuaries Research Institute. https://www.soa.org/globalassets/assets/files/resources/research-report/2026/2026-01-ait184-act-mindset.pdf

Saeri, A. K., Graham, J., Noetel, M., Slattery, P., Ah-king, D., Aittokallio, E., Akindehin, I., Al Mahdi, A., Alhajjar, E., Andersson Lipcsey, R., Ang, G., Azam, C. M., Azaria, A., Balkissoon, R., Barberá, I., Bareato, C., Barry, J., Basehart, M., Bean, A. M., . . . Thompson, N. (2026). Prioritization of risks from artificial intelligence: A Delphi study of 272 international experts. arXiv. https://doi.org/10.48550/arXiv.2606.04490

Slattery, P., Saeri, A. K., Grundy, E. A. C., Graham, J., Noetel, M., Uuk, R., Dao, J., Pour, S., Casper, S., & Thompson, N. (2026). The AI risk repository: A meta-review, database, and taxonomy of risks from artificial intelligence. *Patterns, 7*(5), 101517. https://doi.org/10.1016/j.patter.2026.101517

Appendix A: Methods

Ethics

This project (E-7975) received an exempt determination (Category 2: Educational Testing, Surveys, Interviews or Observation) from the MIT Committee on the Use of Humans as Experimental Subjects (COUHES), in accordance with 45 CFR 46.104(d)(2).

Interview sample

Participants were 13 senior actuarial leaders, a participation rate of 72% of those invited. They were selected to represent a cross-section of institutional contexts, including C-suite executives, chief actuaries, partners, principals, founders, and other senior leadership at insurance carriers, Big Four professional services firms, actuarial technology firms, actuarial consultancies, global reinsurers, academia and professional bodies. The sample spanned a range of actuarial domains and specializations, including life, health, property & casualty, reinsurance, climate and catastrophe risk, regulatory, model risk management, and actuarial professional development. Geographically, about half of participants were based in the United States, with the remainder in Canada, India, South Africa, Australia, and the Asia-Pacific.

Recruitment and interview methodology

Senior actuarial leaders were identified by the SOA Research Institute and put in contact with the MIT research team, who sent email invitations with a brief introduction and interview questions. Semi-structured interviews (60 minutes) were conducted by one or two of the MIT research team on Zoom/Teams, which were recorded and transcribed. Transcriptions were analyzed inductively through data familiarization, coding, and theme development and revision. An iterative and recursive interview approach was used, with themes from earlier interviews guiding later interviews, and participants asked to reflect on existing themes. This was selected given the exploratory nature of the project and to identify areas of consensus and disagreement.

Limitations and other considerations

This study drew on a small, purposive sample of senior actuarial leaders, who were not representative of the profession. By selection, participants are considered more forward-leaning and skew toward the more technical end of the field. While the findings in this study aim to faithfully capture the views of actuarial leaders, these results will be further explored and validated in later phases of the project using survey data and existing empirical research. Prioritization of risks in this phase involved no formal ranking exercise, so the risks reported here cannot be precisely weighted against one another but instead guide further testing.

Finally, recruitment emails and interviews were framed around how AI is reshaping the actuarial profession, which may have primed participants to emphasize change and disruption.

Appendix B: Interview questions

Scene-setting

- *Can you tell me about your area of work and ways that you are seeing AI being used in the last 6 months?
    - Prompt: How is this changing the work that you are doing?
    - Prompt: What is driving some of these changes?

Risks and risk mitigations

- *What are some of the types of risks you are concerned about?
    - Prompt: Examples of incidents
- *What do you think helps in addressing these risks?
    - Prompt: Governance frameworks, organizational policies, work practices, organizational/team culture, codes of conduct
    - Prompt: Where do these approaches fall short?
- [Present MIT risk expert consensus slide and briefly summarize findings] What are your thoughts on these risks? Do you agree with them?

Reflections on existing themes

- Mention that we also wanted your reflections on some of the emerging themes from the research so far. This will help us understand where there are areas of consensus and disagreement.

Non-obvious insights and blind spots

- Is there anything else happening in your work that you think is being missed by people?
- *What is the most important thing you want the profession to understand?
- Is there anything else you'd like to add that hasn't yet been covered?

*Priority discussion questions

Note: Most interviews roughly followed the structure above, with prioritization of questions (*) when time was limited. Depending on participants' area of expertise, some interviews also focused on modeling and managing risks associated with climate and catastrophic events, as well as discussion of actuarial labor market scenarios.