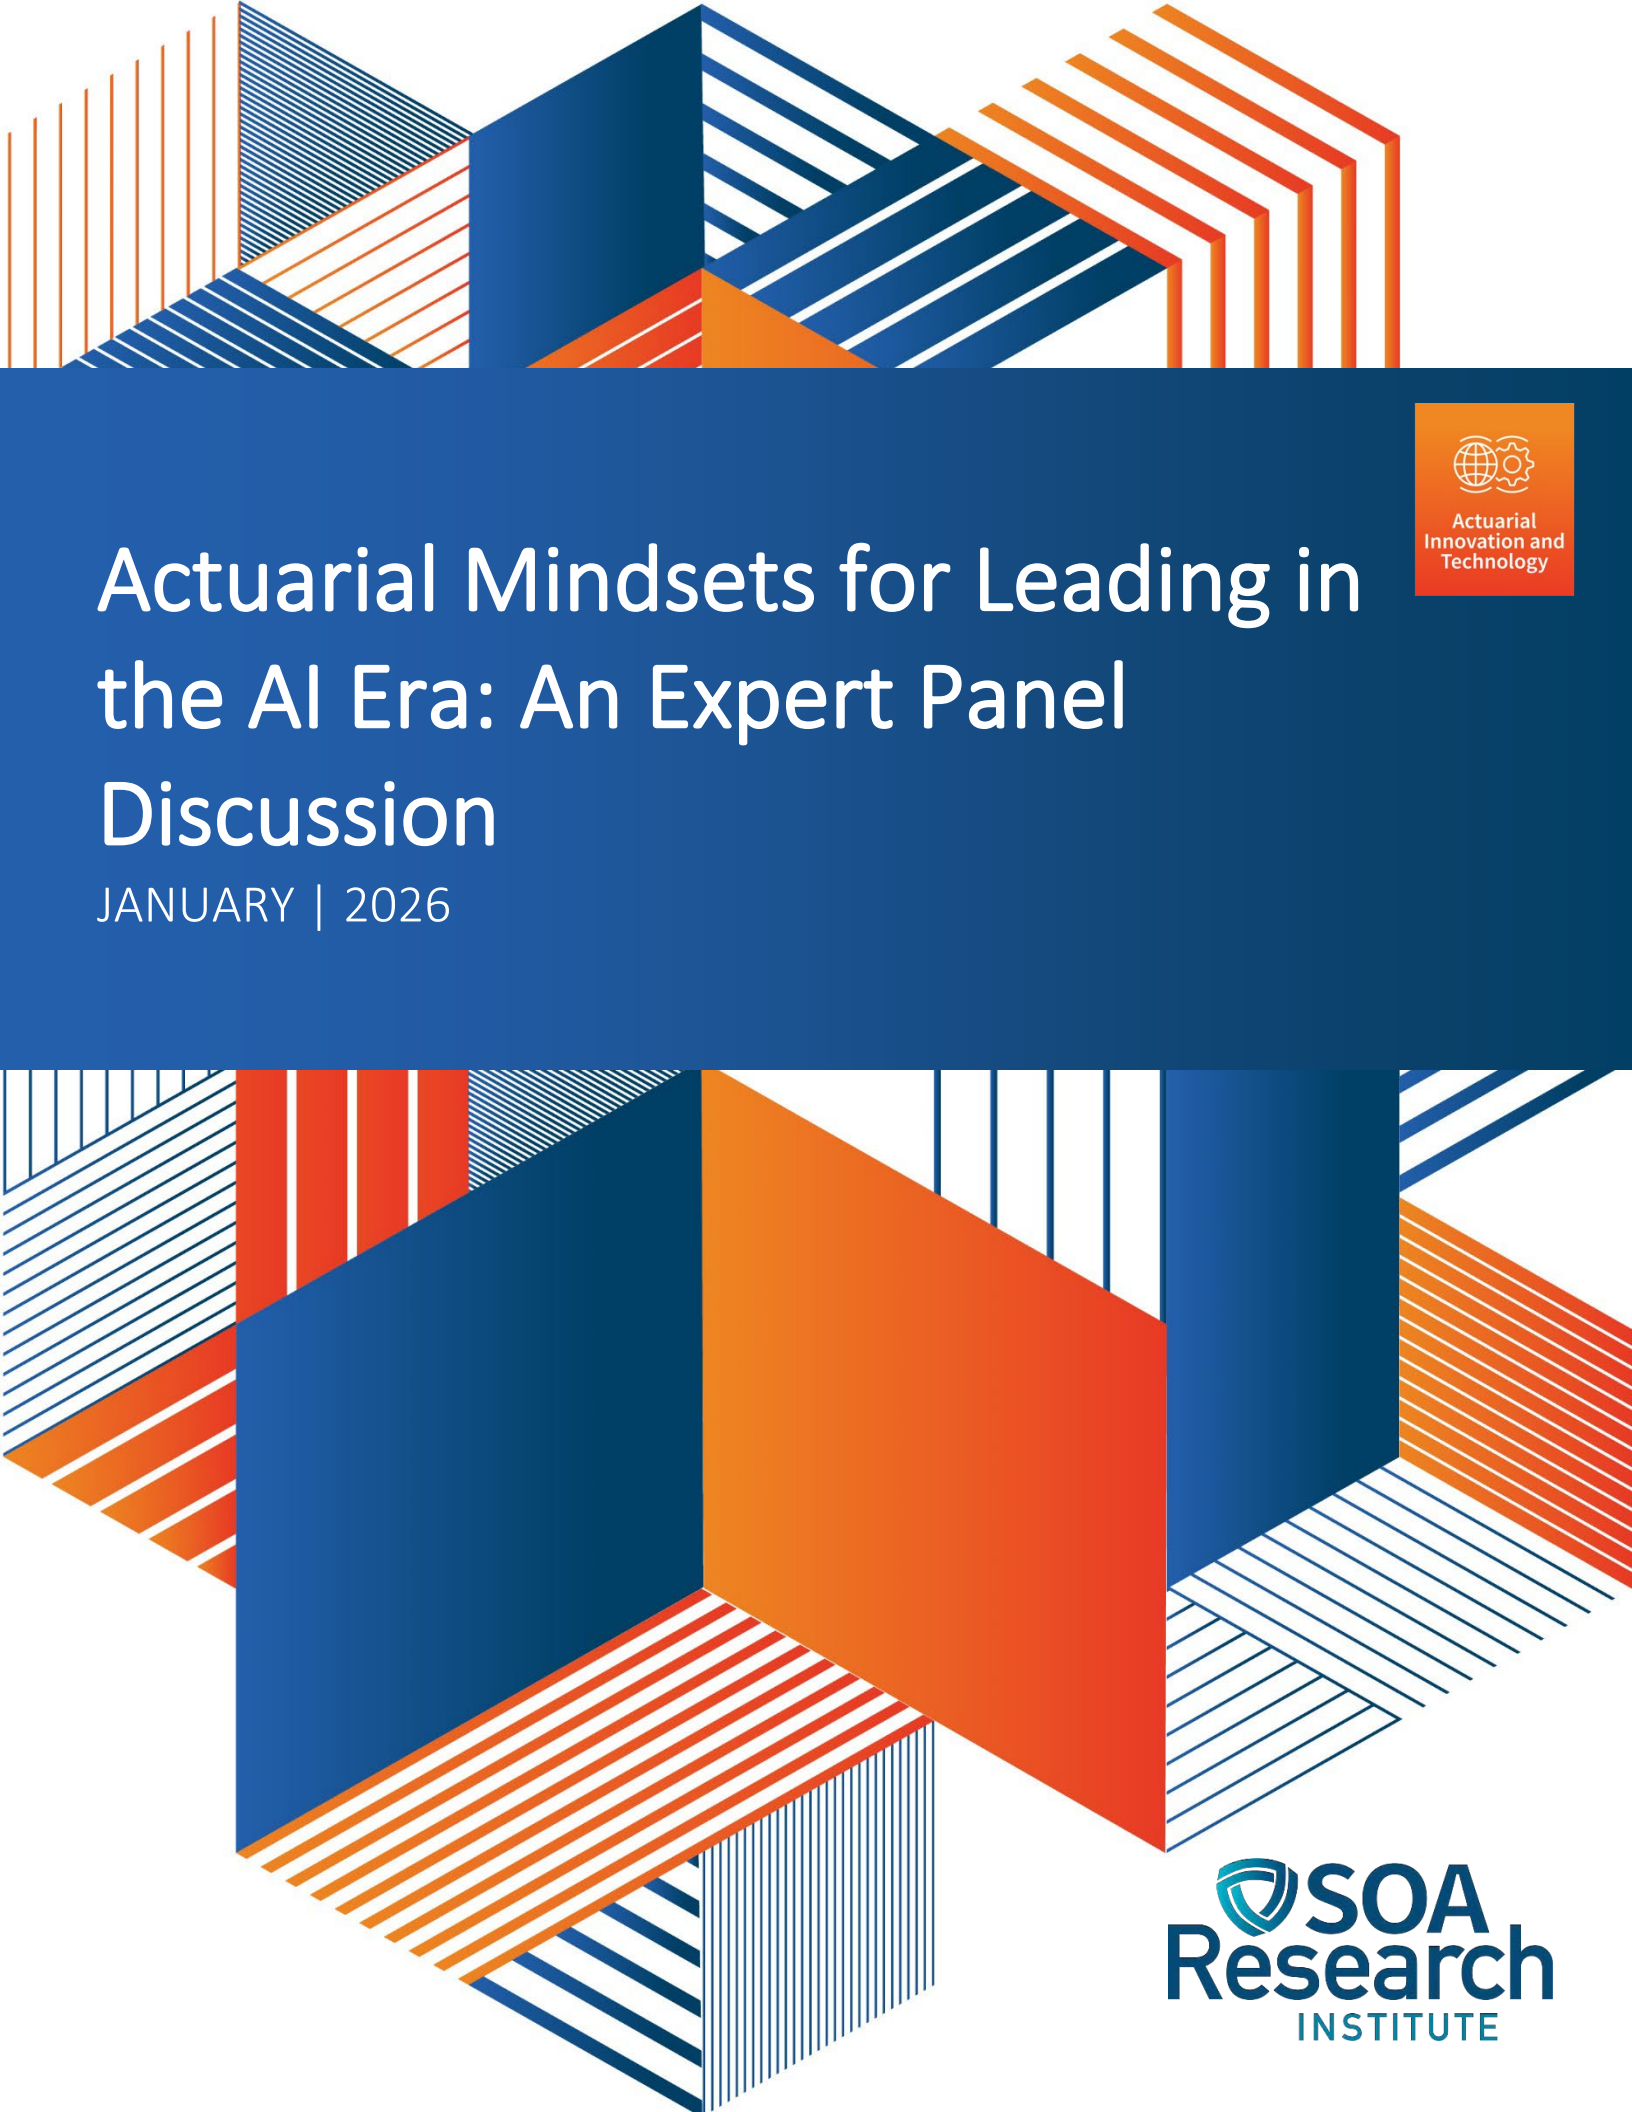

Actuarial Mindsets for Leading in the AI Era: An Expert Panel Discussion

JANUARY | 2026

SOA Research INSTITUTE

Actuarial Mindsets: Leading in the AI Era

An Expert Panel Discussion

AUTHORS Ronald L Poon Affat FSA, FIA C.Act, MAAA, CFA, HIBA

SPONSOR Actuarial Innovation and Technology Strategic Research Program

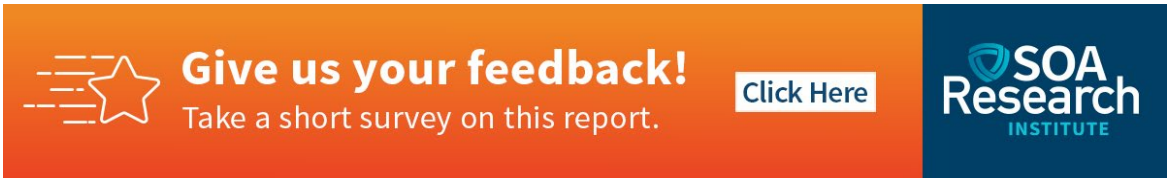

CONTENTS

Actuarial Mindsets for Leading in the AI Era

An Expert Panel Discussion

Executive Summary

This report summarizes insights from a panel of actuarial experts representing academia, insurance, reinsurance, brokerage, and risk consulting. The group was convened to understand at a profession-wide level which enduring mindsets and qualities actuaries will likely need and want to cultivate as Artificial Intelligence (AI) becomes a routine part of actuarial workflows. The discussion focused on identifying the traits that help actuaries adapt to changing tools and evolving analytical environments — qualities such as curiosity, disciplined thinking, the ability to frame problems clearly, and the capacity to integrate new capabilities without losing sight of core professional responsibilities. By drawing on perspectives across practice areas, the panel sought to articulate a forward-looking view of how actuaries can continue to add value as their work is increasingly supported, accelerated, and reshaped by AI.

The work was also motivated by a broader concern: as AI becomes more capable, actuarial decisions continue to affect people directly — claimants, policyholders, and others relying on professional advice. AI cannot interpret these human dimensions on its own, which means actuaries must continue to apply empathy, contextual understanding, and sound judgment when using automated tools. This report therefore examines the capabilities and habits needed to ensure that increasingly technical workflows remain aligned with real-life consequences.

It also examines the capabilities and professional habits shaping actuarial work as advanced analytical systems become embedded in day-to-day practice. The emphasis is on how actuaries think, decide, and communicate in environments where automation can accelerate insight while also introducing new forms of uncertainty.

The discussion further considers how responsibility evolves when established standards develop more slowly than the tools in use, highlighting where core principles remain consistent across practice areas and where specific contexts influence their application.

The panel additionally reflected on how actuarial training and early-career development may need to adapt as increasingly capable systems become widely accessible. Attention was given to how foundational skills are built, how they can be reinforced, and what they may require as the broader environment and toolset continue to evolve.

KEY LEARNINGS

- To maintain professional standards, AI-ready actuaries will maintain accountability, treating AI as informative rather than authoritative. They interpret outputs cautiously, apply judgment, and preserve the ability to explain how and why decisions were made.
- Responsible AI use was described as applying disciplined skepticism without becoming obstructionist, supported by clear documentation, governance awareness, and attention to fairness. These principles apply across all practice areas, regardless of data type or modeling framework.
- Apprenticeship is thought to be most effective when prioritizing the development of deep understanding. Early-career actuaries can benefit from opportunities to cultivate professional judgment, explore models, and learn to challenge outputs constructively, helping them avoid over-reliance on automated answers that lack underlying reasoning.

- Humility is believed to be a core professional trait in the AI era. This includes recognizing the limits of our own expertise, staying open to learning from adjacent fields, and acknowledging when peer perspectives or external insights are needed to refine their thinking.
- Peer review is considered to retain central importance as tools evolve. The ability for one actuary to review, interrogate, and understand another actuarial analysis remains essential for quality control, accountability, and professional credibility, and becomes even more critical when AI-generated content enters the workflow.

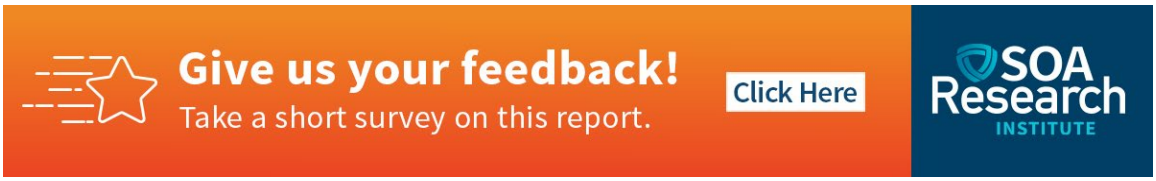

Section 1 What "AI-Ready" Really Means for Actuaries Today

1.1 FRAMING AI-READINESS

The panel began by stepping back from tools and techniques to consider a more fundamental question: what it means for an actuary to be "AI-ready" in today's environment. The opening discussion focused on the non-technical qualities that shape how actuaries engage with emerging technologies — qualities that influence how they approach new systems, adapt to changing analytical conditions, and understand the human expertise they continue to contribute as AI becomes more integrated into their work.

1.2 UNDERSTANDING ONE'S OWN PROFESSIONAL VALUE

The discussion opened by emphasizing the importance of understanding the unique human expertise actuaries bring to human-AI collaboration. This included abilities such as navigating uncertainty, interpreting human behavior, managing complexity, and applying seasoned judgment to real-world systems —capabilities that current AI tools cannot replicate. Participants noted that these activities reflect core professional values and remain essential contributions that only humans can supply. From this perspective, AI-readiness begins with recognizing that effective collaboration with machines still depends on deep human insight, not just technical familiarity with new tools.

1.3 PROFESSIONAL ETHICS AND HUMAN JUDGMENT

Professional ethics and human judgment were described as central to actuarial work in an AI-enabled environment. This ethical duty is not new; it is the bedrock of the profession as outlined in the Society of Actuaries' Code of Professional Conduct. Professional judgment — shaped by experience, ethical reasoning, and contextual understanding — was presented as an essential human function that cannot be replaced by automated systems. This perspective suggested that such judgment is what enables actuaries to evaluate, question, or override machine-generated outputs when needed. Ethical reasoning was also noted as a key factor in weighing consequences, understanding nuance, and ensuring decisions remain aligned with the professional standards codified in the Code, including when AI produces confident but potentially flawed recommendations.

1.4 UNDERSTANDING AI'S ROLES, LIMITS, AND FAILURE MODES

A practical grasp of what AI can and cannot do was described as an important part of using these systems responsibly. AI-readiness was framed not as building models but as understanding how they behave and recognizing when human input is still necessary. This included being able to identify:

- where automated tools tend to perform reliably,
- where outputs may require additional checking, and
- situations that call for review, adjustment, or clarification by an actuary.

Attention was also given to cases where automated results may leave out relevant considerations or simplify issues that require judgment. Several aspects of actuarial work — such as applying context, exercising ethical reasoning, and managing complexity — remain human responsibilities, which support the need for continued involvement alongside automated systems.

1.5 INTELLECTUAL HUMILITY

Intellectual humility was presented as an important mindset when working with AI systems trained on historical data, emphasizing the need to acknowledge the limits of one’s expertise in alignment with the SOA Code of

Professional Conduct requirement to perform services only when qualified.[1] Because these tools may struggle when conditions shift or when new patterns appear, actuaries were encouraged to remain open to the possibility that outputs may not hold under new circumstances, to note when context has changed, and to question whether familiar patterns still apply. This includes staying alert, asking questions, and testing assumptions — an attitude that maintains vigilance rather than assuming automated results will always generalize.

1.6 RECOGNIZING AI AS A TOOL, NOT A PERSON

The discussion noted that people sometimes treat AI as if it were a conversational partner because the systems can sound human, even though they are simply machines generating patterns from data with no awareness or intent. It was suggested that being AI-ready includes keeping this distinction clear, because treating AI like a colleague can lead to misplaced trust or undue reliance on its output. The point was reinforced by outlining the risks that arise when the nature of the tool is misunderstood, including:

- overreliance on automated results,
- trust placed where verification is needed, and
- reduced critical oversight.

Viewing AI as an instrument rather than a decision-maker was presented as a way to maintain appropriate judgment and ensure that responsibility remains with the actuary.

1.7 ABILITY TO THINK CRITICALLY

The ability to "see around the corner" was presented as a key element of critical thinking in an AI-enabled environment. This included identifying what a model may be missing, recognizing which perspectives or dimensions are not reflected in the data, and understanding where risks may arise because the system's view is shaped by its training inputs.

This anticipatory mindset relies on judgment, domain knowledge, and intuition to detect issues that may not appear directly in the model's structure, and actuaries were encouraged to assess outputs through this lens — providing context and insight that automated systems cannot supply.

1.8 KEY TAKEAWAYS

- An AI-ready actuary is believed to start with a clear understanding of the human capabilities they bring to work supported by automated systems. Judgment, ethics, working with uncertainty, and understanding human behavior were described as core elements of actuarial practice. These capabilities shape how actuaries interpret information, recognize what matters in a situation, and bring context to decisions that AI cannot independently assess.
- Readiness is also interpreted as including a practical view of how AI behaves. Participants spoke about knowing where tools tend to perform well, where outputs may require checking, and where human reasoning needs to take precedence. This includes treating AI as an instrument rather than as a collaborator and assessing its outputs with the same scrutiny applied to any other model.

[1] Society of Actuaries, *Code of Professional Conduct* (Effective January 1, 2001), accessed December 18, 2025, https://www.soa.org/about/governance/about-code-of-professional-conduct/.

- The discussion also noted the role of humility and steady critical review. Actuaries were encouraged to revisit assumptions, consider context, and question results that do not align with the underlying situation. This combination of caution, curiosity, and scrutiny supports the ability to identify gaps, anticipate areas a model may not capture, and use automated results in a responsible and informed way.

Section 2 Actuaries Navigating AI Responsibility

2.1 ACCOUNTABILITY IN AN AI-AUGMENTED ROLE

The panel talked about what responsible actuarial practice looks like when AI becomes part of day-to-day work. Their discussion focused on how actuaries understand the tools they use, the principles that guide appropriate oversight, and the professional accountability that remains with the actuary when automated systems inform the analysis. This framing set the stage for examining how responsibility is exercised in an environment where models are more complex, standards are still evolving, and human judgment continues to anchor actuarial decisions.

2.2 KNOWING WHICH AI YOU'RE USING

Responsible use of AI was described as beginning with clarity about the type of system being applied, since "AI is not one thing," and the term is often used as if it referred to a single capability. Actuarial work already engages with several distinct categories of AI, such as:

- large language models (LLMs),
- traditional machine-learning techniques used for pricing and related assumptions,
- deep-learning models applied to unstructured inputs like call transcripts or document intake, and
- time-series or scenario generators used in asset-liability management (ALM) and stress testing.

Understanding which of these categories is in play gives actuaries a more accurate view of each tool's boundaries, strengths, and weaknesses. This conceptual grounding was presented as the starting point for responsible application, separate from any regulatory considerations.

2.3 HUMAN-IN-THE-LOOP GOVERNANCE AS A CORE PRINCIPLE

Human-in-the-loop oversight appeared throughout the dialogue as a recurring point of focus and was described as a foundational element of responsible use. This responsibility was linked directly to established standards such as Actuarial Standards of Practice (ASOP) 56,[2] with emphasis on anchoring work in principles that pre-date current AI capabilities. Within this framing, the panel articulated several operational priorities:

- Intentional model design, where the intended purpose is explicitly understood
- Assessment of potentially poor or adverse outcomes, which must be anticipated in advance
- Continuous validation and real-time monitoring, replacing the older idea of "run a model once and be done"
- Avoidance of model drift, requiring ongoing surveillance
- Rigorous documentation, described pointedly as "documentation, documentation, documentation"
- Justification for reliance on a model and an explicit understanding of when to reject model output

Together, these elements create a cycle of design, monitoring, documenting, and reviewing in which the actuary remains responsible for determining when model output can be accepted or must be set aside.

[2] Actuarial Standards Board, *Actuarial Standard of Practice No. 56: Modeling*, last revised December 2019, accessed December 18, 2025, https://www.actuarialstandardsboard.org/asops/modeling-3/.

2.4 GOVERNANCE EXPECTATIONS FOR AI MODELS

AI was described as another model to handle with the same actuarial rigor applied to valuation or pricing frameworks. This includes:

- Knowing why the model is being used
- Understanding its limits
- Being aware of the data supporting it

The dialogue also pointed to maintaining established governance practices, noting that responsible use is not something fundamentally different, but an extension of existing actuarial discipline applied to a newer class of tools.

2.5 PRACTICING AMIDST EVOLVING STANDARDS

Concerns were raised about applying evolving or incomplete standards, noting that existing professional frameworks may not fully anticipate modern AI tools and that decisions must therefore be grounded in fundamental actuarial principles and broader ethical commitments. A participant observed that actuaries may need to help shape new standards because the current landscape does not yet cover the full complexity of AI-related use cases, and this responsibility involves reflective professional judgment and collaboration with peers as actuaries collectively navigate a rapidly changing environment.

The exchange emphasized that evolving standards do not absolve actuaries from responsibility; instead, they require actuaries to rely more heavily on foundational principles.

2.6 COLLABORATION AS AN ELEMENT OF RESPONSIBLE USE

Participants highlighted that actuaries must collaborate with one another to understand AI's risks and nuances, expressed in the view that "we're all trying to figure this out together," which underscored that responsible use is not an isolated activity but a shared professional effort. This collaborative perspective reinforced the point that effective practice requires a community of practitioners rather than reliance on individual competence alone.

2.7 AI AS A TOOL, NOT A DECISION-MAKER

A participant noted that concerns about AI operating independently are premature, describing it as "another tool in your toolbox," which implies that the actuary retains responsibility for integrating information from AI with other inputs, including human expertise and non-AI analytical tools. The exchange cautioned against treating AI as "its own framework" because current systems do not yet meet that threshold, and responsible use maintains the stance that the actuary — not the model — remains accountable. This perspective treats AI as something powerful but not autonomous, requiring the same judgment and skepticism applied to any other analytic technique.

2.8 OPERATIONAL READINESS AND SKILLED PERSONNEL

The discussion raised concerns about model operations ("model ops") and governance competency, noting the need for both:

- A strong or emerging model ops capability to support lifecycle governance of AI processes and
- Personnel who are genuinely able to design and maintain governance frameworks.

A specific risk was also identified: organizations sometimes assign AI-related responsibilities to individuals without appropriate training or education, which was described as leading to problems and showing that responsible use depends not only on the tools but also on the competence and preparedness of the people overseeing them.

2.9 SUBJECT TO PEER REVIEW

The discussion emphasized the importance of producing work that another actuary could review, even when explicit standards have not been written for the situation, reflecting the Society of Actuaries' Code of Professional Conduct requirement that actuarial communication be clear, appropriate, and understandable to peers.[1] Using the example of reviewing another actuary's work — sometimes in friendly circumstances and sometimes not — the discussion underscored that responsible work must be reviewable, must enable another actuary to determine whether it is "good" or whether "there are problems with it," and cannot rely on an actuary saying "I relied on the AI." It reinforced that responsibility lies with the actuary signing the work, not with the model, and warned that AI-generated output may be handed to actuaries who are expected to sign off without deferring accountability to the system. Overall, this conversation grounded responsible use in accountability, transparency, and professional reviewability.

2.10 KEY TAKEAWAYS

- Responsible use of AI in actuarial practice is viewed as beginning with a clear understanding of the specific systems being applied — whether large language models, traditional machine-learning techniques, deep-learning models, or scenario generators — so actuaries can recognize their boundaries, strengths, and weaknesses and anchor their use in intentional model design, anticipation of adverse outcomes, ongoing validation, monitoring for drift, and disciplined documentation.
- Actuaries are encouraged to treat AI as another model within the traditional actuarial toolkit, knowing why it is being used, understanding its limits, being aware of the data supporting it, and applying foundational actuarial principles where modern standards are incomplete, while retaining human judgment and responsibility rather than deferring decisions to the system and working collaboratively with peers as they navigate emerging risks and nuances.
- Operational readiness and professional reviewability are emphasized as essential conditions, requiring competent model-ops capabilities, personnel capable of designing governance frameworks, and work products that another actuary can meaningfully review to determine whether they are sound or whether problems exist, without relying on the explanation that a model produced the result.

Section 3 AI Fundamentals: Do They Differ by Practice Area?

3.1 CROSS-DOMAIN PRINCIPLES

The panel next considered whether AI fundamentals change across different actuarial domains. Rather than assuming each area requires its own approach, the question was framed around what truly counts as "fundamental" and how those principles apply within the distinct data and modeling environments of Life, Health, P&C, Reinsurance, and Risk Management.

3.2 SHARED FOUNDATIONS ACROSS PRACTICE AREAS

A participant noted that they view the principles outlined earlier in the session — such as judgment, accountability, and careful interpretation of model outputs —as profession-wide expectations rather than domain-specific requirements, a point reinforced by another contribution focusing on professional conduct. This perspective emphasized that responsibilities such as acting with integrity, maintaining competence, serving the public interest, producing robust documentation, conducting validation, and ensuring outputs are "unbiased and fair" are foundational elements of AI readiness that apply uniformly across practice areas.

A further contribution stressed that actuaries, regardless of practice area, are ultimately pursuing similar objectives: improving efficiency, enabling faster decisions, and uncovering novel insights, all within governed processes that uphold the profession's commitment to structure, oversight, and disciplined model use.

Taken together, these perspectives underscored that while data environments and modeling approaches differ across practice areas, the overarching goals and the governance expectations required to achieve them responsibly remain consistent across the actuarial landscape.

3.3 NUANCES BY PRACTICE AREA

Although the fundamental principles were described as broadly shared, participants also pointed to meaningful contextual differences across practice areas. These differences were not presented as changes to core responsibilities but as practical variations that shape how those responsibilities are carried out.

- Life and Health
  This work involves "a lot more personal data," which creates heightened sensitivity around how information is used and requires careful handling of outputs because of the potential impact on individuals.
- Property & Casualty (P&C)
  P&C work relies heavily on structured, data-intensive processes, drawing on long-standing machine-learning and traditional pricing models, and requires substantial data-management effort to ensure that the selected models are appropriate for the volume, quality, and characteristics of the underlying information.
- Reinsurance
  Reinsurance was described as a "pretty wide area" with varied applications, making the choice of model and its alignment with the intended use especially important. The discussion linked reinsurance directly to the need for selecting models that are well suited to the specific context.
- Risk Management
  Risk management often operates "in the areas where there's less data, " requiring actuaries to work with probability distributions developed from judgment rather than observations, making it a domain where expertise and professional judgment play a central role.

Taken together, these points distinguish reinsurance and risk management from data-rich P&C settings and illustrate how AI-related responsibilities must adapt to settings where judgment, rather than extensive data, is the primary driver.

3.4 KEY TAKEAWAYS

- Across Life, Health, P&C, Reinsurance, and Risk Management, actuaries share the same foundational responsibilities — integrity, competence, fairness, sound documentation, disciplined validation, and awareness of model limits — while pursuing common objectives such as improving efficiency, enabling faster decisions, and generating new insights through governed, well-controlled processes.
- The differences across practice areas reflect context rather than principle: Life and Health involve sensitive personal data that require heightened care; P&C relies on structured datasets and long-standing pricing models that demand substantial data-management discipline; Reinsurance encompasses a wide range of applications where selecting models that fit the specific context is essential; and Risk Management frequently depends on judgment-based probability distributions when data is limited. These contrasts illustrate how universal professional responsibilities are applied within distinct environments.

Section 4 Professional Development for AI-Era Judgment

4.1 RETHINKING ACTUARIAL MENTORSHIP

Training for the next generation of actuaries emerged as another central theme in the discussion. Rather than focusing on whether AI might one day automate early-career tasks, the panel turned to how apprenticeship, foundational skill-building, and first-principles understanding should adapt so new professionals develop judgment and depth in an AI-enabled environment.

4.2 UPDATING TRAINING WHILE MAINTAINING ESSENTIAL SKILLS

The conversation opened with the view that AI has not yet automated enough of the entry-level actuarial workload to justify concerns about overreliance. One participant remarked that they “would be surprised” if the challenges faced at that stage were already “so automatable.” This comment underscored that individuals beginning their actuarial careers cannot simply rely on LLMs, because AI is not yet capable of performing the full suite of foundational tasks required in early actuarial work. The group suggested that the real issue, for now, is premature anxiety about this possibility rather than a genuine risk that early-career actuaries will be unable to perform core responsibilities.

It was noted that if AI eventually becomes capable of performing certain entry-level tasks, the profession should not preserve obsolete methods — likened to “using a slide rule” and to the historical period when commutation functions remained on actuarial exams long after they had become impractical. Participants pointed out that early actuarial work once depended on paper-and-pencil commutation functions, skills that disappeared as computerization expanded and entry-level computational jobs vanished, yet the profession continued to grow because new tools increased the value actuaries could provide. This historical perspective underscored a broader principle: when tools evolve, training evolves with them. The implication for AI is similar — while some entry-level tasks may change or shrink, the panel expects the apprenticeship model to adapt rather than diminish, moving away from outdated techniques and toward judgment, oversight, and conceptual depth, which remain essential regardless of the toolset.

4.3 STRENGTHENING EARLY-CAREER FOUNDATIONS

A panel contribution drew attention to existing deficiencies in actuarial training, quite separate from AI. According to this perspective, some practicing actuaries fail to understand first principles; they "run data through actuarial processes and models" without understanding what they are actually doing. Two issues were highlighted:

- Technology asymmetry, in which early-career professionals often have more advanced technology "outside of the company space" than they find inside it, creating expectations — sometimes unrealistic — about what corporate tools can do.
- Weak foundational habits, where some actuaries rely on processes without comprehension, a weakness that "needs to be addressed even before AI." The panel sees cultivating this deep understanding as not just a practical need but a fulfillment of the actuary's duty to maintain competence, as required by the profession's Code of Professional Conduct.[1]

The panelists assert that apprenticeship redesign, therefore, must tackle these pre-existing structural issues, not just those introduced by LLMs.

Building on this, the panel began to explore concrete ways to strengthen foundational development, with one participant saying explicitly that "the Excel actuary kind of needs to come back," arguing that those entering the profession need "time and breathing room" to understand what they are doing. This perspective emphasized slow

thinking, manual exploration, and first-principles work rather than immediately turning to an LLM for fast answers, and it was presented not as nostalgia for outdated tools but as a pedagogical argument that hands-on work strengthens intuition, numerical reasoning, and comprehension in a way that instant answers cannot.

4.4 DEVELOPING FIRST-PRINCIPLES UNDERSTANDING AND MODEL INTUITION

A panelist described apprenticeship as the process of turning mathematically strong students into "professional stewards of risk," noting that the goal of training has not changed even though the toolset has. The concern raised was that modern systems allow students to produce "polished answers that may be right" without "doing the thinking that makes those answers defensible," a pattern intensified when LLMs generate convincing output without supplying reasoning or justification. To counter this, the speaker called for strengthening several elements of apprenticeship:

- Keeping the first-principles foundation strong so actuaries understand how a model works beneath the surface
- Developing the habit of interrogating model behavior, exploring boundary cases, and actively testing where outputs may fail
- Recalling past work with a senior actuary who had an "amazing ability to take any model output and find places where she could break the model.

These points supported the broader view that successful apprenticeship cultivates active, critical engagement with models rather than passive acceptance of polished output.

4.5 MODEL OUTPUTS AS INPUTS, NOT DECISIONS

Another theme emerged: the importance of those entering the profession learning that model outputs are inputs to decision-making — not decisions themselves. One participant stressed their view that actuaries must remain capable of overriding model output and that this responsibility has "always been true," but is "no more true than now."

This recurring point was tied directly to the panelists' strong desire for early-career actuaries to develop sound judgment, contextual awareness, and responsibility for final decisions, so that apprenticeship develops judgment, contextual awareness, and accountability for final decisions rather than allowing tools to dominate thinking — aligning with professional expectations around explainability and reinforcing the view that apprenticeship must cultivate independent judgment, not tool dependency.

4.6 LLMS AS PRODUCTIVITY TOOLS, NOT TEACHING TOOLS

A panelist with data-science teaching experience explained that LLMs are "a very good productivity tool, but a very bad teaching tool," a distinction reinforced by another participant who noted that LLMs excel in helping experts do more, faster, but not in helping newcomers learn. The panel observed that while experts can use them to accelerate tasks and those entering the profession can use them to enhance exploration, they must not become a substitute for learning underlying methods because this risks skipping essential conceptual learning. Effective use of LLMs in an apprenticeship, therefore, depends on active skepticism — asking why the tool is being used, what could go wrong, and how its output has been tested and stressed.

To extend this point, the discussion ended with a metaphor comparing tools in a toolkit: those learning the profession must understand not simply how to "grab a tool and use it for everything," but when to use a tool, when not to use it, how to justify its use, and how to ensure it aligns with standards. This reinforced the perspective that

apprenticeship should cultivate sound judgment about when and how tools apply, not just technical familiarity with the tools themselves.

4.7 KEY TAKEAWAYS

- Across the discussion, participants agreed that AI has not yet automated enough of the entry-level actuarial workload to displace foundational training, reinforcing the need for early-career actuaries to build durable skills while avoiding outdated methods that no longer serve modern practice.
- The panel highlighted longstanding structural gaps in actuarial preparation, such as technology asymmetry and weak first-principles habits, and emphasized that strengthening intuition, manual reasoning, and model-interrogation skills remains essential, particularly as apprenticeships evolve away from obsolete tasks and toward deeper judgment, oversight, and conceptual understanding.
- Several contributors stressed the view that early-career actuaries must learn to evaluate tools rather than defer to them, developing the ability to override models when needed, apply skepticism to LLM-generated output, and understand when and why a tool should or should not be used, ensuring that productivity gains do not replace the development of sound professional judgment.

Section 5 The Learning Journey That Shapes an AI-Ready Actuary

5.1 AI READINESS BUILT OVER TIME

Attention then turned to the learning process behind becoming AI-ready — how actuaries grow into that capability by progressing from early trials with new tools to a steadier, more informed understanding of AI's role in their work, building on long-standing exposure to early predictive modeling techniques such as generalized linear models, random forests, and other traditional machine-learning approaches. The discussion emphasized the learning journey that supports this progression, shaped by accumulated experience, reflection, and increasing fluency with both the strengths and limits of AI.

5.2 USING AI AS GUIDANCE, NOT DIRECTION

AI readiness was described as the ability to ask LLMs useful questions, especially at moments of uncertainty such as "what should I do next?," to draw on AI for guidance rather than literal instructions, and to stay sufficiently critical to recognize when the model's suggestions should be set aside.

This reflects a mindset in which AI supports the thinking process without replacing it. This includes asking LLMs productive questions at uncertain decision points, using AI for directional hints rather than step-by-step instructions, and maintaining the judgment needed to disregard outputs that appear unreliable or misleading.

5.3 HAVING A BALANCED VIEW OF AI'S CAPABILITIES

One participant distinguished between two groups: individuals who use AI in a casual, search-like way and experience small conveniences, and individuals who "truly get it," who reach a clear recognition that AI will change the work environment in significant and rapid ways, along with a corresponding sense of responsibility to help "migrate the company" or organization toward that future. The discussion also connected AI readiness to the Dunning-Kruger effect — a well-established cognitive pattern in which people with limited exposure to a new capability initially overestimate their understanding, only later developing a more calibrated view.

In this framing, early AI users often begin with the impression that AI can do far more than it actually can, but those who become genuinely AI-ready move past this initial stage and recognize that AI has limits, that polished answers can mask underlying errors, and that they need a more measured understanding of strengths and constraints.

A further perspective highlighted that AI-ready individuals have not only experimented widely with AI but also learned to understand it beyond the surrounding hype, becoming comfortable with uncertainty and imperfection, validating outputs rather than accepting them, and communicating about AI in a measured and realistic way.

They focus on when a tool should be used, when it should not, how to use it responsibly, and how it integrates into real workflows — marking the transition from early enthusiasm to mature, operational competence. Another speaker reinforced this view by describing AI-ready people as those who have used AI "extensively," understand how it works at a functional level, remain comfortable with its imperfections and uncertainty, validate outputs rather than accepting them and demonstrate realism by recognizing both strengths and weaknesses. In both descriptions, AI-ready individuals are defined by this same shift from naïve excitement to a more grounded understanding of how AI should be used, when it should not be used, and how it fits into actual actuarial work.

The final contribution emphasized that AI readiness becomes evident in how someone discusses AI — specifically, when individuals can explain it "comfortably and correctly," communicate its limitations with nuance, discuss appropriate use cases, demonstrate awareness of risks and constraints, and understand the need for responsible deployment. This communicative fluency, characterized by balanced and informed descriptions, reflects the person's learning journey: how they have engaged with or experimented with AI, reflected on their experiences,

adjusted their expectations, and ultimately formed a stable internal model of its purpose, boundaries, and role within the organization.

5.4 KEY TAKEAWAYS

The panel views AI-ready actuaries as those who use LLMs as a source of guidance rather than direction, asking productive questions at moments of uncertainty, using models for hints instead of instructions, and maintaining the discipline to disregard outputs that are unreliable or misleading. This reflects a mindset in which AI supports professional reasoning but never substitutes for it.

Panelists also believe readiness involves developing a mature, experience-based view of AI's capabilities and limits, moving beyond early overestimation, recognizing that polished answers can hide underlying errors, and understanding that meaningful change is underway while still keeping a measured grasp of when and how to use AI effectively in practice.

Section 6 Concluding Insights

6.1 RECOGNIZING SKILL GAPS AND LEARNING FROM NON-ACTUARIES

A key perspective underscored that actuaries must recognize they do not have all the skills needed in a future state and must be willing to learn many of those capabilities from non-actuaries, anchored in the explicit statement that actuaries "don't have all the skills we need in a future state" and that many of the needed capabilities will come from outside the profession. This view highlights humility, openness, and cross-disciplinary learning as core professional skills and suggests that the AI era will blur traditional boundaries and traditional domain boundaries, requiring actuaries to collaborate with data scientists, engineers, and technologists whose expertise fills gaps actuaries cannot bridge alone. From this vantage point, the profession and its identity are seen as expanding rather than contracting as actuaries integrate skills from a broader set of disciplines.

6.2 INTERPRETING AI AS PROBABILISTIC SIGNAL

A different contribution framed the essential skill as accountable judgment under uncertainty, emphasizing that AI-generated outputs should be interpreted as uncertain, probabilistic signals. One participant described one common LLM's responses as "a random distribution of words on a page that happen to make sense," requiring actuaries to decide whether those outputs are accurate or reasonable. The discussion linked this interpretive responsibility to machine-learning outputs, neural-network behavior, and broader AI systems. The core idea is that to maintain their role as decision-makers rather than delegating decisions to models, actuaries must apply judgment across all these technologies

6.3 LEARNING TO WORK ACROSS DISCIPLINES

The final contribution highlighted the growing separation between actuarial modeling and the "classical data science" functions within insurers. Actuaries interacting with data-science teams often feel as though they are discussing similar topics using entirely different professional languages. The essential skill identified was the ability to communicate and collaborate across these domains. The discussion identified the importance of acting as translators between actuarial science and modern computational modeling if actuaries are to lead AI-related work effectively.

Section 7 Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Expert Panel participants and others for their diligent work overseeing questionnaire development, analyzing, and discussing respondent answers, and reviewing and editing this report for accuracy and relevance.

Dave Ingram, CERA, FRM, PRM, FSA, MAAA, Risk Management Actuary, Society of Actuaries Board Member, advocate for AI integration in actuarial science.

Shane Leib, FSA, MAAA, Associate Teaching Professor, University of Notre Dame.

Jason Reed FSA MS MAAA, Senior Director, Advanced Analytics, UnitedHealthCare

William White FIA FNZSA. Director, Modelling AI, Aon

Derek Kueker. FSA, SVP, Chief Operating Officer, RGA

Ronald Poon-Affat (Moderator), FSA, FIA C.Act, MAAA, CFA, HIBA—Independent Board Member, Consultant

At the Society of Actuaries Research Institute:

Korrel Crawford, Senior Research Administrator

R. Dale Hall, FSA, MAAA, CFA, Managing Director of Research

Section 8 Appendix A: Panel Discussion Questions

What makes an actuary truly "AI-ready, " especially in terms of non-technical qualities needed to work with powerful and sometimes-wrong AI systems?

What does responsible use of AI look like in actuarial work — especially when regulation and professional standards are still evolving?

Do the fundamentals of AI readiness differ across practice areas — life, health, P&C, reinsurance, risk management---or are the same principles universal?

How should AI-ready actuaries reshape apprenticeship so juniors do not rely solely on LLMs, but still develop the foundational skills needed for explainability and sound judgment?

In one sentence, what's the single most important skill actuaries must develop for the AI era?

When you meet someone who seems truly "AI-ready," what in their learning journey signals that mindset?

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its strategic research programs: aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of topical research available, including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
8770 W Bryn Mawr Ave, Suite 1000
Chicago, IL 60631
www.SOA.org