ACTUARIAL INTELLIGENCE BULLETIN

SOA Research INSTITUTE

September 2026

Actuarial Intelligence Research Bulletin Board 2
How Actuaries Are Moving Beyond AI Experimentation 3
AI Adoption's Real Pain Points Aren't About the Technology 6
Automating Actuarial Models After They're Built 8
I Built a Full Actuarial Business Website Using AI in 15 Minutes. 9
Insurance for AI Users 11
AI and Actuarial Professionalism 15
Unlocking Dark Data in Claims: The AI Engine Behind Enrichment 16
Confabulation as Actuarial Risk: Five Controls That Work in Combination 20
AI and Actuarial Professional Judgment 22
The Evolution of Insurance Asset Portfolio Optimization 23
Virtual Roundtable (Part 2): AI and the Actuarial Profession 26
"Untangling" Our Wicked Problems 29
Editorial Committee 31
About the Society of Actuaries Research Institute 32

Welcome to the September 2026 edition of the SOA Research Institute AI Bulletin! This bulletin serves as a platform for sharing knowledge and fostering collaboration around artificial intelligence within the actuarial community. Explore articles on strategic initiatives, practical tips, and research advancements, all aimed at empowering actuaries to leverage AI responsibly and effectively.

Caveat and Disclaimer
The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute or the Society of Actuaries or its members. The Society of Actuaries Research Institute makes no representation or warranty to the accuracy of the information.

Copyright © 2026 by the Society of Actuaries Research Institute. All rights reserved.

Actuarial Intelligence Research Bulletin Board

SOA RESEARCH INSTITUTE

Welcome to the Actuarial Intelligence Research Bulletin Board! From practical applications and emerging opportunities to governance, ethics, and risk considerations, this bulletin board highlights selected SOA Research Institute work designed to help actuaries understand how AI is shaping the future of our work. You can find all the SOA Research Institute's work on AI here.

Synthetic Claims, Safer Testing

Actuaries can use AI-generated claims data for training, system testing, model development, validation, and scenario analysis when real data cannot be shared. Use pretrained LLMs for structural realism and customized generative models for statistical fidelity—then validate utility, bias, and privacy, especially for rare records.

When AI Sees Faces

Facial recognition can streamline underwriting and fraud detection, but biased training data can create unequal accuracy. Identifying and mitigating potential bias involves testing results by demographic subgroup, assessing data quality and representativeness, documenting limitations, and weighing the trade-offs of mitigation techniques.

The Future of AI in Underwriting

Drawing on perspectives from underwriting, brokerage and intermediary distribution, data and technology, and reinsurance, an expert panel sees AI's near-term value in organizing fragmented evidence, improving consistency, and freeing underwriters to focus on interpretation. The panel frames value beyond speed, encompassing explainability, customer experience, risk quality, and auditability. Across these perspectives, the emerging model pairs AI-supported evidence curation with human judgment, transparent governance, and clear accountability for final decisions.

AI Research Spotlight

When Claims Data Miss the Mark

Healthcare claims show who received care—not necessarily who needed it—so AI models can look accurate overall while failing rural, low-income, underdiagnosed, or intersectional subgroups. This report gives actuaries a three-stage framework to identify potential input, model, and application bias; test performance by relevant and combined subgroups; and select mitigation methods suited to the specific actuarial use.

Synthetic scenarios show that propensity matching may only partly close gaps, missingness-aware approaches can outperform deletion, and fairness techniques that target the chosen objective outperform those that don't. An accompanying Excel Model Card and Decision Assist Tool help document limitations, compare model variants, and support fit-for-purpose decisions in pricing, risk adjustment, utilization management, and other applications.

Copyright © 2026 Society of Actuaries Research Institute

How Actuaries Are Moving Beyond AI Experimentation

JOE ALAIMO, ASA, ACIA

When the Society of Actuaries Research Institute conducted its first AI survey in the summer of 2025, one question stood above all others: Were actuaries using artificial intelligence, or simply talking about it?

The answer surprised many people. While AI dominated headlines, most actuaries reported spending relatively little time using it. Nearly one-quarter of respondents reported no AI use at all, and many organizations were still evaluating where AI might fit into actuarial practice.

Just one survey later, the picture has changed.

The Spring 2026 AI Survey, with responses from 964 actuaries worldwide, shows that the profession has moved well beyond initial experimentation. Respondents represented a broad cross-section of the profession across multiple practice areas, geographies, and organization types. For analysis, respondents were grouped into those with 10 years or less and more than 10 years of actuarial experience. Approximately 25% of respondents had 10 years or less of experience, while approximately 75% had more than 10 years of experience, providing a strong basis for comparing how AI adoption and attitudes differ across career stages. At the same time, because three-quarters of respondents had more than 10 years of actuarial experience, the survey findings primarily reflect the perspectives of experienced actuaries. Overall, the findings indicate that AI adoption has increased, organizations are investing in enterprise AI platforms, and the conversation has shifted from whether AI should be used to how it can be used responsibly.

AI Adoption Is Accelerating

Perhaps the clearest indication of this change is the decline in respondents who reported not using AI at all.

Among actuaries with ten or fewer years of experience, the proportion reporting no AI use fell from 27% in 2025 to just 9% in 2026. A similar decline occurred among actuaries with more than ten years of experience, dropping from 18% to 10%. At the same time, more respondents reported spending between five and ten hours each week learning about or using AI.

TIME SPENT WEEKLY ON AI/ML TOOLS

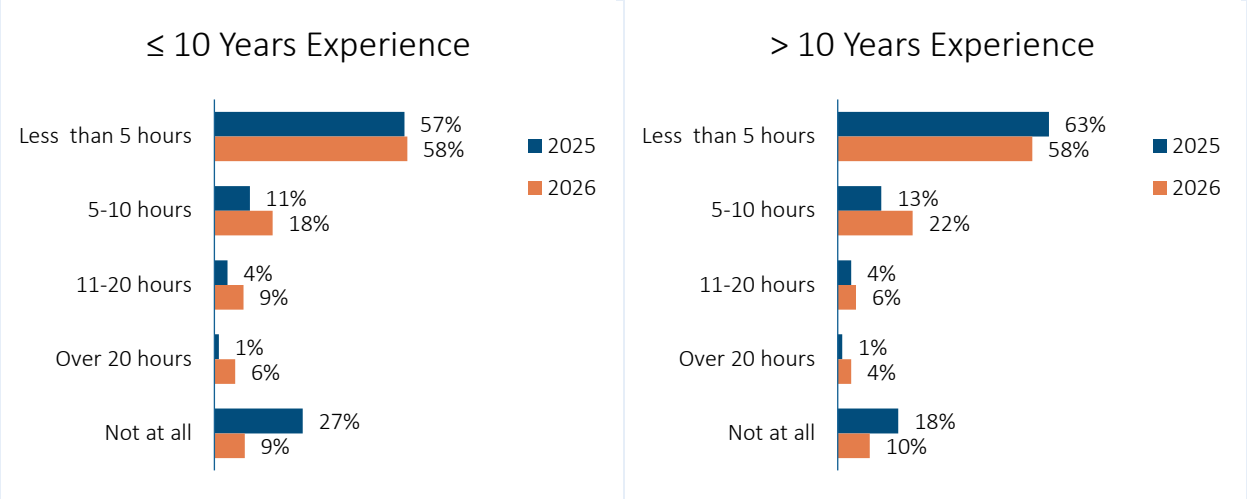

This suggests that AI is becoming part of everyday actuarial work rather than an occasional curiosity.

Copyright © 2026 Society of Actuaries Research Institute

AI Use Is Becoming More Sophisticated

Time savings remain the most frequently cited benefit of AI, reported by approximately four out of five respondents. However, the second survey suggests that actuaries increasingly view AI as more than a productivity tool.

If the first survey documented curiosity, the second documents progress.

Respondents reported growing recognition of AI's ability to improve decision-making, enhance work quality, and expand work product. At the same time, a relatively small proportion indicated that they had not yet realized meaningful benefits, suggesting that while adoption continues to grow, organizations are still at different stages of AI maturity.

AI Is Delivering More Than Productivity

Time savings remain the most frequently cited benefit of AI, reported by approximately four out of five respondents. However, the second survey suggests that actuaries increasingly view AI as more than a productivity tool.

Respondents reported growing recognition of AI's ability to improve decision-making, enhance work quality, and expand work product. At the same time, a relatively small proportion indicated that they had not yet realized meaningful benefits, suggesting that while adoption continues to grow, organizations are still at different stages of AI maturity.

BENEFITS FROM AI

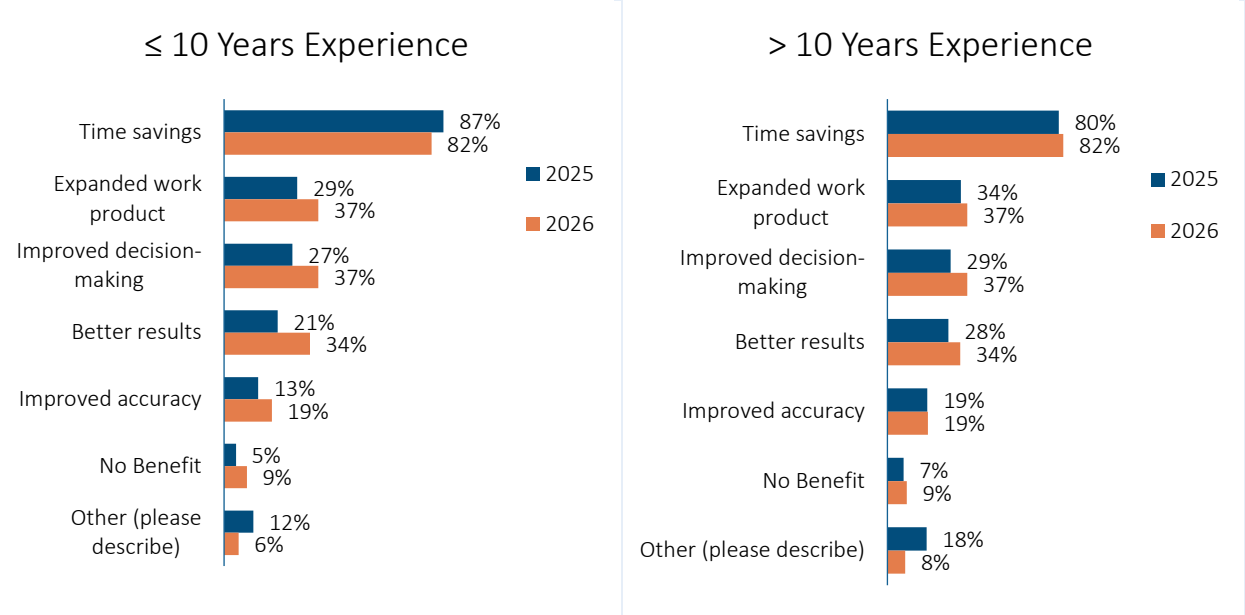

Governance Has Become the Central Conversation

One of the most significant shifts between surveys is not how actuaries use AI, but how they think about it.

In 2025, organizational barriers centered on regulatory compliance, skill gaps, and uncertainty regarding AI's return on investment. Those concerns remain important, but many organizations appear to have moved beyond asking whether AI should be adopted.

Instead, respondents are increasingly focused on governance.

Copyright © 2026 Society of Actuaries Research Institute

New survey questions and response options reveal growing attention to organizational AI policies, actuarial standards, training resources, and governance frameworks. At the same time, "wait and see" approaches to AI governance declined substantially, suggesting that many organizations are now actively developing governance practices rather than simply observing how the technology evolves.

Accuracy Matters More Than Ever

As organizations gain experience with AI, concerns have also evolved.

The survey introduced "Accuracy and Hallucinations" as a specific response option after it emerged repeatedly in the first survey's open-ended comments. It quickly became the most frequently selected governance concern among both early-career and experienced actuaries. Respondents also identified data and model bias, explainability, cybersecurity, and compliance as ongoing priorities.

These findings reflect the actuarial profession's natural focus on validation, accountability, and risk management.

Enterprise AI Is Becoming the Norm

Another notable finding is the rapid emergence of enterprise AI.

Approximately six in ten respondents reported that their organizations now provide enterprise AI models. Microsoft Copilot was by far the most commonly deployed platform, followed by ChatGPT and Claude. Some organizations have even begun deploying proprietary large language models tailored to their own environments.

This suggests that AI is becoming an organizational capability rather than an individual experiment.

Looking Ahead

The first AI survey established a baseline for understanding how actuaries were beginning to engage with AI. The second survey tells a different story.

AI adoption is increasing. Organizations are investing in enterprise tools and formal governance. Actuaries are discovering broader applications for AI while maintaining a strong emphasis on professional judgment, validation, and responsible use.

Perhaps the most encouraging finding is that the profession has not abandoned its traditional strengths as AI adoption has accelerated. Instead, actuaries appear to be approaching AI in the same way they approach any emerging risk: thoughtfully, analytically, and with appropriate governance.

As this longitudinal survey continues, it will provide an increasingly valuable view of how AI evolves within the actuarial profession. If the first survey documented curiosity, the second documents progress. The next surveys may well show how AI becomes an established part of actuarial practice.

Joe Alaimo ASA, ACIA is CEO of Procomp Consulting Inc.

Copyright © 2026 Society of Actuaries Research Institute

AI Adoption's Real Pain Points Aren't About the Technology

MAY ZHENG, FSA, FCIA

Most talk about AI in insurance starts with what the tools can do: draft a note, evaluate a claim, produce some code. But the evidence from 2025 and 2026 points somewhere else. The hard part is rarely the model. It is the training that never happened, the messy data, the rules that keep changing, and the human judgment we count on. Here are five places where adoption gets stuck, and what each one means for insurance companies and their employees.

One point sits under all of them. Almost every real deployment in insurance companies keeps a human in the loop principle. AI helps the underwriter or examiner; it does not replace them. That setup is also how carriers defend their work to regulators. However, the critical point is - it only holds up if the person doing the review is really reviewing, not just signing off.

Pilots That Never Pay Off

Nearly every life carrier now uses AI somewhere. Surveys put it near 87%.[1] But a 2025 MIT study found that about 95% of generative-AI pilots never reached real production value,[2] and people now talk about "pilot fatigue." The problem is usually not about technology. It is the data. About 78% of carriers say theirs is not ready,[3] because old systems were built to store records, not to train models.

Training Gaps that Become Fairness Gaps

Training is the cheapest high-value fix, yet most firms skip it. About 68% of workers using AI received no formal trainings in the past year, and those who are trained could save roughly twice the time compared to those who taught themselves.[4] When training is rare, the people who get good tend to be the ones who already had an edge: younger, better educated, and by some counts more often men, because managers push men to use AI about 23% more than women.[5] In a business that will be checked for unfair treatment of customers, the same gap inside your own staff is worth watching.

Shadow AI

Between 57% and 78% of workers who use AI do it outside any approved process, and 38% to 43% admit pasting sensitive work data into tools without permission.[6] The average AI-related breach costs more than $650,000.[7] Bans

---

[1] LIMRA and LOMA, Equisoft, and Universal Conversion Technologies, "LIMRA and Equisoft Report Reveals 78% of Global Life Insurers Believe Data Readiness Is the Biggest Challenge to Getting Value from AI," January 15, 2025, https://www.equisoft.com/insights/insurance/limra-and-equisoft-report-reveals-78-of-global-life-insurers-believe-data-readiness-is-the-biggest-challenge-to-getting-value-from-ai; Zineb Jazouli, "5 Trends Reshaping Life Insurance in 2026," Equisoft, February 23, 2026, https://www.equisoft.com/insights/insurance/5-trends-reshaping-life-insurance-in-2026: Directional survey figures.

[2] Challapally, A., Pease, C., Chari, P., & Raskar, R. (2025). *The GenAI divide: State of AI in business 2025*. MIT Media Lab, Project NANDA (Networked Agents and Decentralized Architecture). https://nanda.media.mit.edu and IBM (2026).

[3] LIMRA and LOMA, Equisoft, and Universal Conversion Technologies, "LIMRA and Equisoft Report Reveals 78% of Global Life Insurers Believe Data Readiness Is the Biggest Challenge to Getting Value from AI," January 15, 2025, https://www.equisoft.com/insights/insurance/limra-and-equisoft-report-reveals-78-of-global-life-insurers-believe-data-readiness-is-the-biggest-challenge-to-getting-value-from-ai.

[4] Daniel Jolles and Grace Lordan, Bridging the Generational AI Gap: Unlocking Productivity for All Generations (London School of Economics and Political Science, The Inclusion Initiative, and Protiviti, October 2025), https://www.lse.ac.uk/tii/assets/documents/Bridging-the-AI-Generational-Gap.pdf

[5] Mary Noble-Tolla, "New Research: Women Use AI Less Often at Work and Get Less Credit," LeanIn.Org, June 5, 2026, https://leanin.org/articles/findings/ai-women-gender-gap-data/. .

[6] Adam Lawrence, "Is Rising AI Adoption across the US Workforce Creating Shadow AI Risks?" *IBM*, November 3, 2025, https://www.ibm.com/think/insights/rising-ai-adoption-creating-shadow-risks. Microsoft and LinkedIn, "AI at Work Is Here. Now Comes the Hard Part," *2024 Work Trend Index Annual Report*, May 8, 2024, https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part. National Cybersecurity Alliance and CybSafe, *Oh, Behave! The Annual Cybersecurity Attitudes and Behaviors Report 2024–2025* (2024), 88–89, https://20740408.fs1.hubspotusercontent-na1.net/hubfs/20740408/CYBSAFE-Oh%2C%2BBehave%21%2BThe%2BAnnual%2BCybersecurity%2BAttitudes%2Band%2BBehaviors%2BReport%2B2024-2025.pdf. Ozan Ucar, "What Is Shadow AI?," *Keepnet*, updated March 12, 2026, https://keepnetlabs.com/blog/what-is-shadow-ai.

[7] IBM, *Cost of a Data Breach Report 2025* (IBM Corporation, 2025), https://www-api.ibm.com/adobe/assets/urn:aaid:aem:f6d39474-dcd2-4607-a5fd-039674143ac9/original/as/cost-of-a-data-breach-report-2025-executive-summary.pdf.

Copyright © 2026 Society of Actuaries Research Institute

do not stop this; they push it out of sight. In life insurance the stakes are higher because the data is: medical records, finances, beneficiaries, claim files. The fix is not "be careful." It is giving people a safe and helpful tool that is easier to use than their own account.

Speed That Is Not Real Progress

A 2025 study watched skilled developers work in code they knew well. They thought AI made them 24% faster. Afterward they still believed it helped by 20%. But it made them about 19% slower.[8] The lesson carries over. AI helps most on simple tasks that are easy to check and can slow you down on work that takes judgment and must hold up under review. That is actuarial work, and actuarial mistakes are quiet: a reserve can be wrong without looking wrong, and it often surfaces late, in audit or validation.

There is also a skills worry. A highly-quoted study found that heavy AI users scored lower on critical thinking, especially younger ones, but it is a single study and only shows a link, not a cause.[9] The simpler argument does not need it: in a human-in-the-loop model, your value is the judgment you built by doing the work yourself. If your day turns into approving AI output, that skill fades, and the review becomes a rubber stamp.

In a human-in-the-loop model, your value is the judgment you built by doing the work yourself.

A practical rule: use AI freely for low-risk drafting, use it with care when you are testing logic, and keep it as a second check only, never the final answer, where the call really depends on judgment.

Rules Moving Faster than Readiness

Only about a quarter of insurance executives say they are confident they could pass an outside AI governance review in 90 days, and 44% blame governance problems for AI projects that failed.[10] Meanwhile the rules are tightening. The NAIC's evaluation tool is moving into real exams, Colorado already requires testing life models for unfair discrimination, and the EU treats life and health risk scoring as high-risk.[11] The good news is this is not new ground for actuaries. It lines up with the Control Cycle and ASOP 56 Modeling. What is missing is an agreed method for the testing, especially fairness testing, at the level regulators now expect.

What It Adds Up To

None of these problems are solved by a better model, and all of them land on the person who signs the work. Four things help:

- **Train your people.** It is the cheapest fix and the one most company skip.
- **Give them approved tools instead of bans.** A safe, easy tool beats a rule people work around.

---

[8] Shane Leib, "The Actuarial Efficacy Paradox: Why AI Feels Fast but Works Slow," AI Bulletin (Society of Actuaries Research Institute, May 2026), 4–6, https://www.soa.org/globalassets/assets/files/resources/research-report/2026/2026-05-ait170-ai-bulletin.pdf Expected ≈24% faster, believed ≈20% faster, measured ≈19% slower.

[9] Michael Gerlich, "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," Societies 15, no. 1 (2025): article 6, https://doi.org/10.3390/soc15010006 Single study, correlational (n = 666); shows a link, not cause.

[10] Grant Thornton Advisors LLC, 2026 AI Impact Survey: The AI Proof Gap (2026), 4, 7, 9, https://www.grantthornton.com/content/dam/grantthornton/website/assets/content-page-files/advisory/ai-lp/infographic/ai-impact-survey-2026/pdf/grant-thornton-2026-ai-impact-survey.pdf ≈24% very confident of passing an independent governance review in 90 days; 46% tie failures to governance/compliance.

[11] National Association of Insurance Commissioners, Big Data and Artificial Intelligence (H) Working Group, Artificial Intelligence Systems Evaluation: Optional Supplemental Exhibits for State Regulators, draft, version 4.0, accessed September 4, 2026, https://content.naic.org/sites/default/files/inline-files/AI%20Systems%20Evaluation%20Tool%204.0%20%28Clean%29.pdf; Colorado Division of Insurance, "Regulation 10-1-1: Governance and Risk Management Framework Requirements for Life Insurers', Private Passenger Automobile Insurers', and Health Benefit Plan Insurers' Use of External Consumer Data and Information Sources, Algorithms, and Predictive Models," 3 Code of Colorado Regulations 702-10, amended effective October 15, 2025, §§ 2–6, https://www.sos.state.co.us/CCR/GenerateRulePdf.do?fileName=3+CCR+702-10&ruleVersionId=12182&utm; European Parliament and Council of the European Union, Regulation (EU) 2024/1689 of June 13, 2024, Laying Down Harmonised Rules on Artificial Intelligence, Official Journal of the European Union L 2024/1689 (July 12, 2024), annex III, para. 5(c), https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng.

Copyright © 2026 Society of Actuaries Research Institute

- **Measure real results, not how fast it feels.** Track rework and review time across the full job.
- **Protect judgment.** Decide where AI is a second check, not the author, so the person in the loop still matters.

May Zheng, FSA, FCIA is Director at Moody's Analytics in Toronto, Canada.

Automating Actuarial Models After They're Built

MASON ROBERTS, FSA, MAAA, MBA

Every actuary has built something like this. Claims lag triangles pulled into Excel. Multiple methods run side by side to develop completion factors. Large claims pulled out and analyzed separately so they don't distort the pattern. Then a manual review, a selection of final factors, maybe an average PMPM calculated as a cross-check. Multiple files, multiple people, a fair amount of copy and paste, and more than one meeting before anyone signs off. Building that process the first time can take a week or more. Running it every cycle after that is still a multi-week affair.

The only step in that whole process where an actuary is adding value is the judgment call at the end, selecting the final estimate. Everything upstream of that, pulling the data, running the triangles, calculating the factors, is mechanical. It doesn't need a person doing it by hand every single cycle. It needs to be built once, correctly, and then left alone.

That pattern is a candidate for automation. I've spent the last few months moving analytics code at my firm from something a person runs by hand to something that runs itself, mostly on Snowflake. One set of code that used to take a full day to run and check step by step now runs as a scheduled package in under twenty minutes. That leaves more time for careful review and analysis.

Six Things a Script Needs Before It Can Run Itself

Turning a working script into something that runs unattended isn't just wrapping it in a scheduler. It needs to:

1. Produce correct results no matter when or how many times it runs (idempotency)
2. Have its step-by-step dependencies made explicit, not assumed
3. Get its warehouse, database, and credentials from the job definition, not an interactive session
4. Fail loudly instead of silently, with someone notified and a record kept
5. Be scheduled, whether on a calendar or triggered when new data lands
6. Get validated end to end, unattended, before anyone trusts it

The scheduling step is the easy part. The above is where the real work is.

Where AI Adds Value

I let AI write nearly all of my documentation. For me, writing code is enjoyable, writing it up is not. It's also useful for catching bugs and cleaning up code. It once flagged that an adjustment factor, I'd built had the numerator and denominator handled inconsistently across different factors. The math was right since I had applied it correctly, but the approach wasn't consistent, and that made it harder for anyone else to follow. An easy fix, and a clearer script because of it.

Copyright © 2026 Society of Actuaries Research Institute

The tradeoff is that AI is doing a lot of the production, which means you need to be doing more of the review. Documentation and code both need a careful read before they go anywhere, or you are at risk of shipping a well-formatted hallucination.

Platform Considerations

My experience is mostly with Snowflake, but the pattern should apply to all organizations. Databricks has its own job orchestration, Lakeflow Jobs, with the same dependency and scheduling logic. SQL Server shops get there through SQL Server Agent jobs or Change Data Capture paired with a trigger. Companies without a platform-native scheduler may use Apache Airflow. What's universal, regardless of platform, is idempotency, parameters that live outside the code, and failure handling that isn't silent. What's platform-specific is how much of that plumbing the platform builds for you versus what you have to build by hand.

Governance and the Real Pitfalls

Once code runs unattended, build governance around it the same way you'd govern anything else. A drift check just compares this run's results to last runs and flags what changed. A graceful failure depends on how critical the failure is: something serious should halt, notify, and hand over enough detail to diagnose it; something more minor should get flagged and let the process continue. Either way, the goal is the same, an appropriate response paired with enough information to act on it.

The nearest miss I've had wasn't a broken pipeline, it was sending out deliverables that were quietly inconsistent with each other. That's what a good drift and consistency check is for.

The bigger risk isn't technical, its spending real time automating something that was never a good candidate: work that needs heavy customization every time, or that isn't refreshed often enough to justify the build. Before you invest, make sure the people who'll monitor and maintain the pipeline, often a data engineering team, understand the value and have room for it in their own workload.

Where to Start

Clean up the code first. There's no downside to that regardless of what happens next. Pull out the hardcoded values. Then go learn what your organization's systems can do and find the person who's already tried this. Someone probably has, and they can save you the lessons they learned the hard way.

Mason Roberts, FSA, MBA is Chief Actuary at Paramean Solutions.

I Built a Full Actuarial Business Website Using AI in 15 Minutes.

SYED RAZA, FSA, MAAA

It started as a quiet weekend experiment. No grand plan, no team, no development budget. Just curiosity, a browser tab, and an online platform for building, running, and hosting software and apps directly in your web browser.

By the time I made my second cup of coffee, I had a fully functional, professionally designed actuarial marketplace website live on the internet.

Let me tell you how that happened; and more importantly, why it should matter to you.

Copyright © 2026 Society of Actuaries Research Institute

The Idea

A while back, I wrote a LinkedIn post that struck a nerve.[12] The premise was simple: actuaries working inside insurance companies often run into highly technical questions; regulatory grey areas, reinsurance structures, niche pricing problems; and have no obvious person to call. The expertise they need doesn't live in-house. Consultants are expensive and slow to engage. There's no "break glass in case of emergency" option.

So, I floated an idea: what if there were a marketplace where actuaries could book a one-hour advisory call with a vetted expert? Pay for exactly the insight you need, nothing more.

The post got a strong response. But it stayed an idea; until last weekend.

Enter Browser-Based Development

For those unfamiliar, browser-based development environment are available. You don't install anything. You don't configure servers. You open a tab, write or prompt your way through some code, and the tool handles the rest; running, hosting, and deploying your project in the same place. Think of it as handing the engine room over to someone else so you can focus on steering.

I decided to use it to finally bring the "Ask an Actuary" concept to life, not necessarily to launch a business, but to see what was possible.

> If you have a workflow problem and a few hours, you may be closer to a solution than you think.

Here's what I did: I took my original LinkedIn post, dropped it into my LLM, and asked it to generate a full product specification for the website. Then I fed that spec into the browser-based development tool as a prompt.

Five minutes later, I had a working website.

Not a rough mockup. Not a skeleton with placeholder text. A *full-fledged* site; with a clean interface for actuaries seeking advice, expert profiles, service categories spanning pricing, product development, reinsurance, regulation, and more.

Over the next ten minutes, I iterated. A few more prompts gave me an admin panel to approve or reject expert applicants, improved search filters, and refined UX. Fifteen minutes, start to finish.

What This Actually Means

I want to be careful not to reduce this to a "wow, AI is amazing" moment and leave it there. Because the implications run deeper than that, and they're worth sitting with.

Not long ago, turning a business idea into a working website required either significant technical skill, a developer on your payroll, or weeks of back-and-forth with an agency. The barrier wasn't just money; it was time, access, and the translation layer between "what I want" and "what a developer can build."

That barrier is collapsing.

---

[12] Syed Raza, "Actuaries, Have You Ever Wished for a Hotline to Industry Experts?" LinkedIn, May 15, 2025, https://www.linkedin.com/posts/syed-raza-a2589722_actuaries-have-you-ever-wished-for-a-hotline-share-7328785140870975488-4cas/?utm_source=share&utm_medium=member_desktop&rcm=ACoAAATCDf0BddsXVvPnmHcMT0pnJA7KGBR9ohE.

Copyright © 2026 Society of Actuaries Research Institute

What I experienced last weekend isn't a party trick. It's a signal of a broader shift: the gap between having an idea and being able to test it in the real world is narrowing at a pace that most professions; including ours; haven't fully reckoned with.

For actuaries, this matters in at least two ways.

First, the practical: tools like the one I used combined with AI assistants mean that prototyping internal tools, client-facing dashboards, or exploratory models no longer requires a full IT project. If you have a workflow problem and a few hours, you may be closer to a solution than you think.

Second, the strategic: our profession has always derived value from technical complexity. We are trained to work where the problems are hard and the data is messy. But as AI makes *execution* easier, the premium shifts toward judgment, domain expertise, and knowing which questions are worth asking in the first place. That's a space actuaries are well-positioned to occupy, provided we stay close to what's changing.

A Note on Staying Current

I post regularly on LinkedIn about AI and how it's reshaping the actuarial landscape; the tools worth knowing, the capabilities worth experimenting with, and the risks worth understanding. If you're not already following along, I'd encourage you to.

And of course, that's also what this newsletter is for.

The profession is not standing still. Neither should we.

Full Article available at: https://www.theactuaryweekly.com/p/i-built-a-full-actuarial-business-website-in-15-minutes-here-s-what-that-means-for-all-of-us

Syed Raza, FSA, MAAA is Founder & Principal Actuarial 360 Solutions

Insurance for AI Users

KETAKI BORKAR

For *Part 2 of a two-part North America AI risk insurance series*

The insurance market for artificial intelligence remains in a period of rapid transition, particularly for organizations that use AI rather than develop or sell AI products themselves. Historically, many AI-related losses have been absorbed within traditional insurance policies that do not refer explicitly to artificial intelligence, an approach often described as "silent AI" coverage. Cyber, professional liability, crime, employment practices, property, fiduciary and other conventional policies may therefore respond to an AI-related loss if the underlying cause of action falls within existing policy wording.

That position, however, is becoming less reliable. Insurers are increasingly introducing affirmative AI wording through endorsements, standalone products and specifically defined coverage, while some are simultaneously adding exclusions that narrow or remove coverage for losses involving AI. For insureds, the practical consequence is

Copyright © 2026 Society of Actuaries Research Institute

that a claim which appears to be a conventional cyber, crime, E&O or employment practices loss may respond differently once AI becomes part of the facts.

The issue for AI users is therefore not necessarily the emergence of entirely new categories of loss, but the interaction between new technology and insurance contracts designed for risks that predate it.

Cyber with "Affirmative AI" Endorsements

Coverage: Affirmative

This is traditional cyber coverage that has been expressly broadened to address AI-related incidents involving external models, AI workloads, prompts, vector stores, or cloud-based AI dependencies. The policy response may depend on whether the use of an external AI platform is contemplated within the wording.

Example: An employee uploads confidential client files into an external generative AI tool that is not properly restricted, leading to unauthorized disclosure and privacy claims.

Professional Liability / E&O (Service Firms)

Coverage: Silent

This responds when professional firms use AI tools in service delivery and clients allege that AI-driven mistakes led to negligent advice, missed standards, or financial harm.

Example: A law firm relies on AI-assisted research and files a brief containing fabricated case citations. The client alleges negligent representation after the matter is delayed and sanctioned.

Commercial Crime

Coverage: Silent

This is increasingly relevant where AI enables deepfakes, voice cloning, synthetic identities, and sophisticated social engineering fraud directed at finance and treasury functions.

Example: Accounts payable receive spoofed vendor instructions supported by deepfake video verification and updates payment details, sending legitimate invoices to a criminal-controlled bank account.

Employment Practices Liability (EPLI)

Coverage: Silent

This addresses workplace claims arising from the use of AI in recruiting, screening, monitoring, promotion, discipline, or termination decisions.

Examples:

- An employer uses an AI résumé-screening tool that disproportionately filters out candidates from a protected class. Applicants bring a discrimination claim after the pattern is uncovered.
- A company relies heavily on productivity-monitoring AI to discipline and terminate employees. A former employee alleges wrongful termination and unfair bias in the decision process.

CGL / Property / Business Interruption

Coverage: Silent

This becomes important when AI controls real-world operations and a failure causes bodily injury, property damage, or an operational shutdown at the insured's own site.

Copyright © 2026 Society of Actuaries Research Institute

Examples:

- An AI-controlled conveyor routing system sends products into a jam that damages machinery and halts a distribution center for two days, creating repair costs and lost income.
- A building-management AI mismanages temperature settings during a cold snap, leading to frozen pipes, water damage, and partial suspension of operations.

IP-Specific Insurance (User Defense)

Coverage: Affirmative

This is targeted coverage for enterprises that use generative AI outputs in commercial products, campaigns, or codebases and want defense support for resulting IP disputes.

Examples:

- A marketing agency launches a campaign built on AI-generated images and later receives a copyright claim alleging the visuals are substantially similar to a protected work.
- A software company deploys code generated with an AI assistant into a customer-facing product and is later accused of incorporating infringing or improperly licensed code.

Personal Cyber / Identity Protection

Coverage: Silent

This is consumer-oriented coverage for individuals and households targeted by AI-enabled fraud, identity theft, extortion, and account takeover schemes.

Examples:

- A family is targeted with a deepfake kidnapping scam using a cloned child's voice, and they incur emergency response, counseling, and financial loss before the fraud is discovered.
- A criminal uses AI-enhanced identity fabrication to open accounts and take over existing profiles, leaving the insured to manage restoration, legal, and cleanup costs.

Fiduciary Liability

Coverage: Silent

This protects those overseeing employee benefit plans when AI tools are used in plan administration, selection support, or participant-facing guidance and that use is later challenged.

Example:

- A benefits committee adopts an AI-supported recommendation engine for retirement plan allocations. Participants later allege the tool steered them toward imprudent options without adequate oversight.

Contingent Business Interruption (CBI)

Coverage: Silent

This responds when the insured's income is disrupted because a critical third-party AI or cloud provider experiences an outage or service degradation.

Examples:

- A company's customer-support platform depends on an external LLM provider. A prolonged upstream outage prevents service delivery to paying clients and causes measurable revenue loss.

Copyright © 2026 Society of Actuaries Research Institute

- A startup's underwriting assistant runs entirely on a third-party AI API. A multi-day disruption forces manual fallback, missed transactions, and contractual service credits.

Financial Institution Professional Liability (FIPL)

Coverage: Silent

This is tailored to financial institutions using AI in lending, underwriting, advisory, or other regulated decision-making contexts where errors can trigger civil and regulatory exposure.

Example:

- A lender uses AI-based credit scoring and is accused of producing disparate impact against protected borrowers, leading to class action allegations and regulatory scrutiny.

Standalone Regulatory Investigation & Defense

Coverage: Affirmative

This is meant for organizations that expect regulator attention over AI usage even before a formal lawsuit exists, especially in hiring, healthcare, finance, and consumer-facing deployments.

Examples:

- A consumer-facing company receives a civil investigative demand regarding whether its AI chatbot made misleading disclosures or unfairly handled customer interactions.
- A state regulator opens an inquiry into an employer's use of AI in hiring and promotion decisions, requiring extensive document production, legal review, and model-governance analysis.

Conclusion

For AI users, the central insurance challenge is increasingly one of coverage interpretation and ensuring risk is explicitly covered rather than simple product availability. A deepfake fraud, biased hiring algorithm, AI-assisted professional error or third-party model outage may fall within established insurance architecture, may require affirmative AI wording, or may be restricted by an emerging AI exclusion.

Policy review is therefore becoming as important as policy purchase. Organizations need to identify where AI influences customer advice, employment decisions, regulated activities, operational systems and third-party dependencies, and then examine whether existing policies, endorsements, sub-limits and exclusions reflect those exposures. As insurers move from implicit treatment of AI toward more deliberate underwriting positions, relying on silent AI coverage without examining the wording is becoming progressively less defensible.

Ketaki Borkar is Senior Director, North America P&C, Sapiens.

Copyright © 2026 Society of Actuaries Research Institute

AI and Actuarial Professionalism

JUAN HERRERA, FSA

Actuarial professionalism has always been central to the profession. The need to explain results and document procedures so another actuary could reach the same conclusion from the same information is just as important today as it was forty years ago, when I first entered the profession.

Some version of the "black box" has existed for forty years. The first actuarial modeling software I encountered for PCs appeared in the mid-1980s to support cash flow testing. When the Actuarial Opinion and Memorandum Regulation (AOMR) was adopted by the NAIC in the early 1990s—with timing varying by state—actuaries issuing Actuarial Opinions were required to perform cash flow testing.

In the 1980's actuarial modeling system developed by Shane Chalk and his company, was built for cash flow testing while the practice was still voluntary. The software was complex, required substantial data, and often took a day or more to produce results. It was originally written in APL, short for A Programming Language, whose strength was executing compact commands that conserved computing power. That same efficiency, however, made it difficult to understand exactly how calculations were performed. The opining actuary would then use the results to determine whether the organization passed the required cash flow tests.

The same questions remain relevant today. Are the model's results accurate? Was the input data reliable? How should the work be communicated? The Actuarial Standards of Practice (ASOPs) identify what actuaries need to consider and provide guidance for performing their work. As AI becomes more of a black box, it can offer even less visibility into how information is transformed.

Two years ago, the American Academy of Actuaries' Committee on Professional Responsibility published an article on generative AI and professionalism. I spoke with Kathy Wong, a member of the committee who helped develop the paper.

> *Juan: What AI-related risks concern actuaries today that were either less apparent or less developed when the article was written?*
>
> Kathy: AI development continues to evolve, in terms of both the pace at which it is developing and the complexity of its capabilities. Thus, all the concerns raised in the Committee on Professional Responsibility discussion paper continue to be relevant and probably even more critical to consider. However, one risk that I see as more important today is that AI has become so integrated into how many actuaries do their work that they start seeing it as just a normal work tool, like Excel. In the paper, we focused on what actuaries need to think about if they are going to employ AI in an assignment - getting agreement from the principal to use it, vetting the results, making sure they are qualified to use AI. Today I think the usage of AI has become so much more common and normalized that actuaries often don't consider even mentioning that they have used it, don't think about asking permission, and don't think about the risks involved.
>
> *Juan: Do you believe actuaries now need a baseline level of AI literacy to meet professional competency requirements? If so, what knowledge should every actuary possess?*
>
> Kathy: I don't think that it is necessary that all actuaries have some level of AI literacy, if they recognize "what they don't know." AI is a tool. It's becoming more important to understand it and to use it properly and to recognize the potential pitfalls when you see something produced by AI.
>
> But think about when spreadsheets and Excel first came into use. Some actuaries working at that time became experts and could use that knowledge to their advantage. Other actuaries learned the basics, which allowed them to understand others' spreadsheets and perform simple tasks. And some actuaries of

Copyright © 2026 Society of Actuaries Research Institute

that generation never learned Excel at all—but they could still function competently using the more primitive tools that they knew. Professional competency would become an issue if an actuary tried to use AI without having that basic understanding or literacy, or if an actuary were working on or reviewing a project where someone provided AI-generated information and the actuary did not understand what it meant.

Juan: What professionalism challenges remain unresolved?

Kathy: I'm not sure I'd use the term "unresolved." Professionalism needs to underlie any work an actuary is doing, and the basic precepts of our Code of Professional Conduct have stood the test of time. In other words, professionalism considerations with respect to AI are embedded in the Code as written. Where the issues can arise is in the application of the Code to this newer AI era. The point of our paper was to show how the precepts of the Code could apply to AI and what actuaries might do from a practical standpoint to be sure they are exhibiting professionalism.

Using the term "unresolved" more broadly, there is certainly room to expound further on the practical applications of the Code, related to the more advanced and more prevalent AI usage that now exists.

Juan: *How should actuarial standards evolve?*

Kathy: Talking about standards gets us more into the nitty-gritty aspects of actuarial work, and we could expect individuals or organizations to provide their own viewpoints and examples to the Actuarial Standards Board for consideration. Myself, I can immediately see two areas in which actuarial standards *might* need to evolve. One is the idea of updating existing standards in light of specific effects AI is having on the area of work covered by the standard. Should AI-specific wording be added for various tasks? Could that even rise to the level of there being a section in the standard for "Use of AI in . . . (reserving, or pricing, or whatever the subject is)?

The other question is whether AI usage should have its own standard of practice. Are there macro-level considerations that should be spelled out that apply to any AI usage? I personally lean toward AI-specific wording being added to existing standards rather than development of a separate AI standard.

To download the professionalism's paper on "Generative AI and Professionalism" you can visit the Academy's website at https://actuary.org/ and navigate to their discussion papers under professionalism. You can also download the paper at this link https://actuary.org/wp-content/uploads/2024/10/professionalism-paper-generative-ai.pdf

Juan Herrera, FSA is a retired Health Actuary and a volunteer with the SOA.

Unlocking Dark Data in Claims: The AI Engine Behind Enrichment

JESS KANG, FSA, CERA, FSAS

The ideas, framework, and practitioner observations in this article are grounded in my own experience across APAC insurance data and analytics work. I used Claude (Anthropic's AI Assistant) as a writing collaborator to help me

Copyright © 2026 Society of Actuaries Research Institute

structure, articulate and refine the script. All arguments and conclusions are my own, and I have reviewed the final text in full.

Every insurer's claims data warehouse looks tidy from the outside. The schema is clean, the fields are well-defined, and the values are numeric or categorical: cause code, reserve, diagnosis code, sum insured. What that schema hides is how little of what was known about a claim ever makes it into a model. The surveyor report, the medical file naming a specific doctor and a course of treatment, the police report identifying a third party, the first-notification call: all of it is collected, all of it is paid for, and almost none of it is structured. This is the dark data of claims, and for most of actuarial history, illuminating it has been too expensive to attempt at scale. Artificial Intelligence (AI) is what has changed that calculation.

The Dark Data Problem & The Extraction Architecture

The core AI job in claims enrichment is turning unstructured evidence into structured, model-ready features. That work splits into two stages, and both have been transformed by recent AI progress. The first stage is extraction, converting a scanned document, a photograph of a form, or a block of free text into machine-readable content. Layout-aware document parsers now handle the messy reality of insurance paperwork like rotated scans, handwritten notes, faxed-and-refaxed PDFs, far better than classical optical character recognition. The second stage is interpretation, pulling specific, schema-defined fields out of that text. This is where large language models have done the most to change what is operationally feasible. A task that once required a bespoke natural-language-processing pipeline, built and tuned separately for every document type, can now be prototyped in days with a well-designed extraction schema and a general-purpose model. Fields like prescribing doctor, clinic address, diagnosis, medications and dosages, diagnostic tests, treatment dates, damage category, root-cause commentary and salvage assessment fields, previously locked inside PDFs, can now be queried almost as soon as the document lands in the system.

Copyright © 2026 Society of Actuaries Research Institute

THE FIVE-TIER REFERENCE ARCHITECTURE

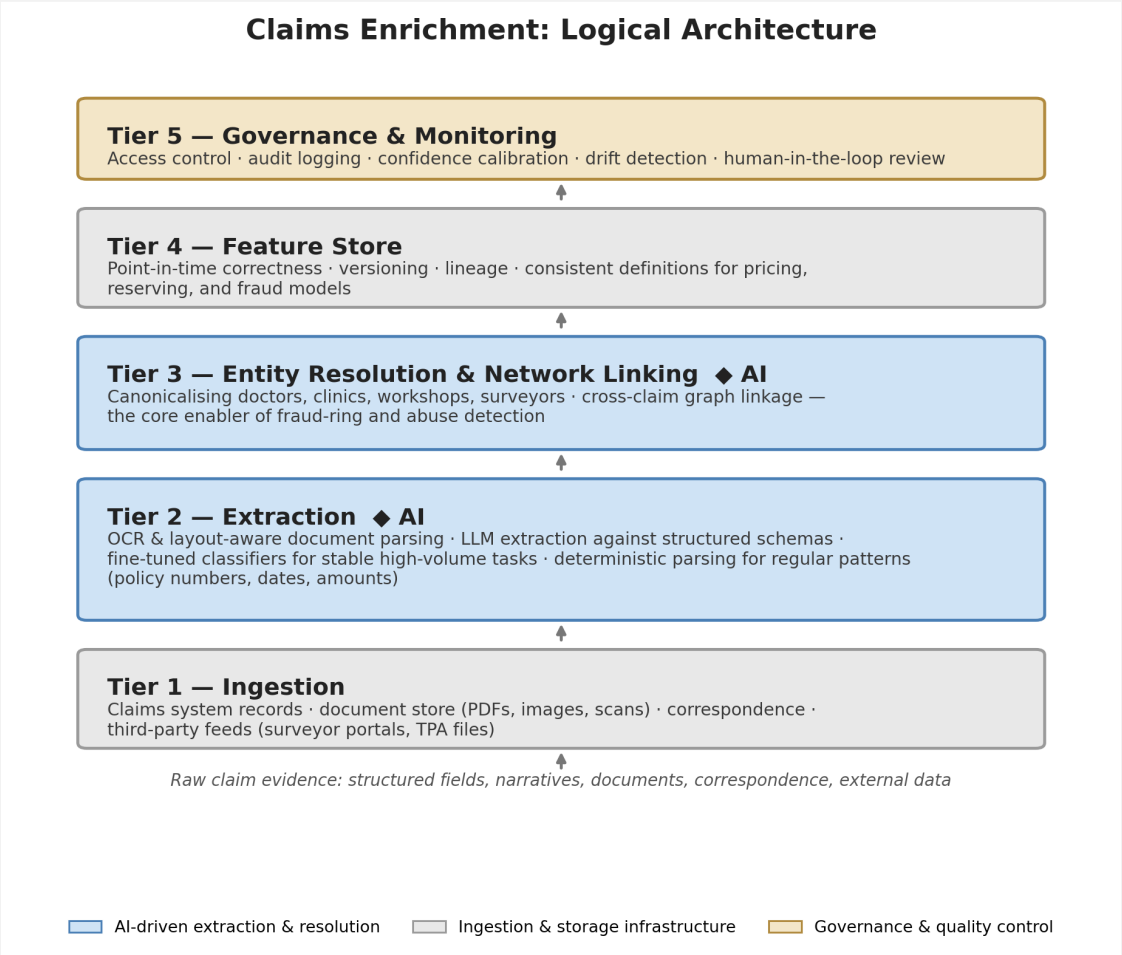

Source: author with ChatGPT.

AI does most of its work in Tiers 2 and 3, extracting fields from documents and resolving them into linked entities, inside a structure which ingestion, storage, and governance layers keep that output auditable.

Not every extraction task should go to a general-purpose model, and this is a point worth making precisely because it is easy to over-apply AI once it works well on the hard cases. High-volume, stable tasks, such as normalizing a cause code or classifying a document type, are often better served by a small, fine-tuned classifier. Highly regular patterns, such as policy numbers, dates, and dollar amounts, are frequently more reliable when pulled out with deterministic, rule-based parsing than with a model at all. The craft of a good enrichment architecture is routing each extraction task to the cheapest, most reliable tool that can do it, and reserving large models for the genuinely unstructured, context-dependent work they are suited to.

Unlocking Fraud Networks & Cross-Claim Intelligence

The single highest-leverage application of this AI capability may be leaning towards entity resolution and network analysis, and it is worth explaining why. Conventional fraud scoring works on claim-level features like velocity, amount, time to notification, and catches individual opportunistic fraud reasonably well. It is nearly blind to organized fraud rings, because ring activity is defined by relationships across claims like the same clinic, the same treating doctor, the same repair workshop, the same surveyor, appearing repeatedly, and sometimes even under slightly different spellings and addresses. Once AI extracts those names and addresses from medical reports,

Copyright © 2026 Society of Actuaries Research Institute

invoices, and surveyor documents, a resolution layer canonicalizes them into single entities, and a graph structure links claims through them for downstream analysis. A clinic appearing on fifteen unrelated soft-tissue claims within three months, all with similar medication patterns, becomes visible for the first time. That signature simply does not exist in a structured claims record; AI is what makes it detectable at all, and it is the same underlying capability which is structured extraction plus resolution that also sharpens pricing segmentation, surfaces early severity signals for reserving, and identifies recoverable third parties for subrogation.

Calibration, Governance & The Actuarial Boundary

None of this works, however, if the output of the AI layer is treated as if it carried the same certainty as a keyed-in claim amount. An extracted diagnosis is a model's best estimate, conditional on document quality, parsing quality, prompt design, and model quality, and it should be labelled as such. Every enriched attribute needs a calibrated confidence score, one where a model reporting 90% confidence is correct about 90% of the time, and explicit thresholds for auto-acceptance, human review, and rejection. Low-confidence extractions should route to a human reviewer rather than pollute a feature store silently. And because enrichment pipelines are typically built on third-party models, they inherit those models' upgrade cycles, a version change can quietly shift extraction behavior across millions of historical records. Without version-pinning and before/after comparison on any model upgrade, an enrichment pipeline can invalidate the very point-in-time correctness that actuarial use demands.

That discipline of deciding where AI's probabilistic outputs meet the actuarial bar for provenance, auditability, and reproducibility is squarely an actuarial responsibility, not a data-engineering afterthought.

This also sets a firm boundary on where AI-derived features belong in the decision chain. Enriched features are excellent at triage: flagging a claim for review, ranking subrogation files by recovery potential, surfacing a possible fraud-ring linkage for an investigator to examine. They should not, on their own, drive irreversible adverse decisions about a customer like a claim denial or a fraud referral to law enforcement because residual extraction error persists even at high confidence. The right mental model is that AI decides which claims a human looks at first, not which claims a human never sees.

The economics of this capability are also moving targets. Model pricing and performance are improving quickly enough that a use case profitable today may look different within a year, which argues for continuous review rather than one-off business cases. But the fundamental shift is durable, AI has taken claims data enrichment from a theoretical aspiration to an operationally realistic program. The insurers that build the extraction, resolution, and governance discipline to use it credibly will have pricing, reserving, fraud, and subrogation capabilities their peers cannot match on structured data alone. That discipline of deciding where AI's probabilistic outputs meet the actuarial bar for provenance, auditability, and reproducibility is squarely an actuarial responsibility, not a data-engineering afterthought.

This article is a condensed excerpt. For the full picture, including the broader enrichment framework, additional use cases across pricing and reserving, and a deeper look at governance and data quality, see the complete article "Unlocking Dark Data in Claims" at LinkedIn.

Jess Kang, FSA, CERA, FSAS, is a senior data and analytics leader in the APAC insurance industry.

Copyright © 2026 Society of Actuaries Research Institute

Confabulation as Actuarial Risk: Five Controls That Work in Combination

MARTIN MELCONIAN

The single most dangerous failure mode for actuaries using generative AI is not prompt injection. It is not data leakage. It is not even bias. It is confabulation.

Confabulation is NIST's technical name for what most people call hallucination: the production of confidently stated but factually wrong content. NIST AI 600-1, the *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, published on 26 July 2024, defines it precisely as erroneous or false content, including outputs that diverge from the prompt or contradict previously generated statements in the same context. It happens because large language models are trained to produce plausible next tokens, not true ones. Plausibility and truth are correlated but not identical. Where they diverge, the LLM produces confabulation.

Why Confabulation Is the Dominant Actuarial Risk

For actuarial work this is the dominant risk because the cost of being confidently wrong is enormous and the failure mode is hard to detect. An LLM that confabulates an ASOP section number, a regulation paragraph, a numerical assumption, or a citation looks exactly like an LLM that does not. The text is grammatically correct, professionally toned, internally consistent, and false.

The American Academy of Actuaries' 2024 discussion paper *Actuarial Professionalism Considerations for Generative AI* is direct. Where GenAI supports actuarial conclusions, ASOP 56 applies in full. ASOP 23 applies to the inputs. ASOP 41 applies to the outputs. The Code of Professional Conduct applies. APS X2 v1.1 (effective 30 January 2026) governs peer review. Incidental autocomplete that does not touch actuarial conclusions sits outside this scope. But once an LLM is producing research synthesis, drafts for external communication, or code that goes into a validation pipeline, the full professional stack engages.

Five Controls that Work in Combination

Five mitigations work, and a competent actuarial GenAI deployment uses several in combination. None eliminates confabulation completely. Together they reduce it to a level that, combined with appropriate human review, is consistent with professional standards.

- Retrieval-Augmented Generation (RAG). The architecture in which the LLM is given access to a defined corpus of authoritative documents and required to ground answers in retrieved passages. The corpus is processed into a vector database. When the user asks a question, the system retrieves relevant passages, supplies them as context, and instructs the model to answer using only those passages. For actuarial use, a well-constructed corpus might contain every relevant ASOP, TAS, and APS, current NAIC bulletins and state insurance department circulars, the NIST AI RMF Playbook and GenAI Profile, and the firm's internal methodology. The corpus must be version-controlled. RAG does not eliminate confabulation, because the LLM can still ignore, conflate, or interpolate between retrieved passages. It dramatically reduces the rate and gives every output a defined provenance the actuary can verify.

- Structured output. Modern LLM APIs (OpenAI's structured outputs, Anthropic's tool use, Google's function calling) allow developers to constrain output to a JSON schema. A properly designed schema can force the model to produce, for every claim, a citation to a source span. For every numerical assertion, the source field and value. For every recommendation, a confidence level and caveats list. Schema-conformant output can be programmatically validated. Every required field present. Every numerical value within plausible range. Every citation pointing to a real document. This is the automated first line of defence against confabulated output.

Copyright © 2026 Society of Actuaries Research Institute

- **Citation requirements.** Every factual claim the LLM produces is accompanied by a pointer to its source, identifying document, section, and where possible sentence or character span. The system prompt should make this explicit. Every factual claim must be supported by a citation. If no supporting citation exists in the retrieved passages, the model says so rather than producing the claim. This transforms the actuary's review task from "is this true?" to "does the citation support this?" A confabulated claim either has no citation, has a citation that does not support it, or has a citation to a document or span that does not exist. All three surface directly.

- **Evaluation suites against ground truth.** The artefact category most actuarial teams have never built and that is most worth investing in. An eval suite is a structured collection of test cases for which the correct answer is known, run against the system at every change and at regular intervals, with results tracked over time. For an actuarial RAG assistant on regulatory citations, test cases take the form "given this draft paragraph, produce the relevant ASOP and TAS citations." Expected outputs are created by a human actuary in advance and stored in version control. Aggregate recall, precision, F1, hallucination rate, and source faithfulness are tracked across model versions. Eval suites give the team a quantitative measure of system quality, catch regressions when providers update underlying models without notice, and produce MEASURE evidence mapping directly onto MEASURE 1.1, 2.1, 2.3, and 2.5. Twenty to fifty test cases are enough to start. A mature suite has several hundred.

- **Human review gates with sign-off semantics.** The most fundamental control. For any LLM output that enters a work product or supports an actuarial conclusion, explicit human review and sign-off, recorded, is the discipline that satisfies the standards. The sign-off is not a tick-box. It is the actuary asserting professional responsibility for the content with the same weight as asserting responsibility for any other actuarial work. The specific semantics calibrate to use case. Research synthesis the actuary consumes directly needs lighter-touch citation verification. Drafting work consumed externally gets the same standard of review as any other deliverable. Code generation gets peer review and test execution. Embedded production systems make the system itself the work product, and sign-off is on the system as a whole.

Combined with version pinning, so provider updates go through the AI Governance Committee rather than being silently absorbed, these five controls produce a defensible confabulation stack. The Code of Professional Conduct and the Actuaries' Code are clear. "The model told me" is not a defence. The five controls are what make professional responsibility discharge-able in practice.

The full article with the fully worked example of a regulatory citation assistant (architecture, Govern and Map artefacts, eval thresholds, Manage runbook) and the reusable LLM usage policy template is in Part 4 of the NIST AI RMF for Actuaries series on the Globebyte Insights site at https://globebyte.com/insights/nist-generative-ai-profile-for-actuaries.

Martin Melconian, Founder and CEO of Globebyte

Copyright © 2026 Society of Actuaries Research Institute

AI and Actuarial Professional Judgment

NAVEEN AGARWAL, MBA

Insurance industry is fast moving to a distribution, capital management, and everything in between model. The middleware of underwriting, claims, servicing and operations is being rapidly replaced by AI agents. It's clear the core differentiation over time, will live in distribution and capital management. Actuaries will play a differentiating role as this unfolds. Yet, one of the most anxious questions I get from CFOs is—"Will AI replace actuaries?". It's a reasonable ask since no analytical job is AI proof.

In actuarial context, a single axis organizes most of the answer: Does the future resemble the past? Are you merely seeking sophisticated interpolation? Is the data good, and is the underlying process stable? Can you trust the analysis or is it made up to please a certain point of view? Is answer dependent on the question asked or is the real challenge, "what are the questions?". Actuarial work spans that whole range—which is why the discipline of *how* an actuary reaches for AI matters as much as *where* it fits. The fact of the matter is, AI does not erase human judgment; it relocates and concentrates it.

What AI Can Substitute

The computational core of actuarial analysis is increasingly automatable. Data ingestion and cleaning, experience-study mechanics, triangle projection, GLM and gradient-boosted model fitting, valuation-system refactoring, first-draft documentation, and regulatory-text summarization are all tasks where AI is faster and often more consistent than a human. In P&C personal lines this has been true for years—gradient boosting is now standard in pricing pipelines. The marginal hour an actuary once spent hand-fitting a frequency model is largely gone, and that is a good thing.

Where AI Complements

Between full automation and irreducible judgment sits the largest and most interesting territory. Here AI proposes and the actuary disposes. In P&C, models surface candidate rating variables and interactions from telematics and unstructured claims text that a human would never be able to synthesize. In Health, predictive models stratify populations for care management and sharpen risk adjustment. In Life, pattern detection can flag anomalies in mortality experience or stress a stochastic economic-scenario far more thoroughly than a manual review. The actuary's job shifts from generating the analysis to interrogating it: is the signal real or is it leakage, is the variable permissible, does the pattern have a mechanism or is it an artifact of the sample?

Where Judgment Is Irreducible

Three things AI cannot do for the profession, and they matter most where the stakes are highest.

1. Reasoning through structural breaks. AI is weakest precisely when the future stops resembling the past—and each practice area has its own version of this. Life actuaries set longevity-improvement assumptions that pay out over decades; no amount of historical fitting tells you whether mortality improvement continues, plateaus, or reverses. Health actuaries face therapy and utilization shocks—GLP-1s, cell and gene therapies, a post-pandemic care backlog—that have no analogue in the training data. P&C reserving actuaries must judge social inflation, litigation-funding trends, and a non-stationary climate that undermines the stationarity cat models assume. In every case the data runs out exactly where the money is, and someone must decide what happens beyond it.

2. Operating inside constraints that are legal and ethical rather than statistical. The most predictive model is frequently not the permissible one. Proxy discrimination, unfair-discrimination statutes, rate-filing and justification requirements, ACA rating rules, and explainability expectations from regulators all mean the actuary must choose a model that is defensible, not merely accurate. That is a judgment about fairness and

Copyright © 2026 Society of Actuaries Research Institute

the public interest, and increasingly it extends to governing the AI itself—validating it, documenting its limitations, and standing behind its use.

3. **Accountability. An actuary signs an opinion.** A model does not. Professional standards, the duty to the public, and personal answerability to boards and regulators attach to a named human who can explain the reasoning and own the consequences. This is not a technical gap that AI will close with more scale or better architecture; it is a feature of what a profession is. And most vital.

The Shape of the Role

The net effect is that the actuarial role hollows out in the middle and thickens at both ends. Upstream, framing the right question and governing assumptions becomes more valuable. Downstream, communicating uncertainty to non-technical decision-makers and owning the sign-off becomes more valuable. The computation between them gets commoditized.

The real risk for actuaries is not redundancy but complacency—the deskilling that comes from rubber-stamping model output no one fully understands. An actuary who cannot reconstruct why a number is what it is, and cannot defend it, and cannot catch it when it is wrong—someone who has given control to AI of thinking not doing. The skill and willingness to be answerable to their analysis.

AI will do more of the actuarial "middle" work every year. What it will not do is carry the deep understanding of the analysis. That's the signature. That remains the actuary's—and increasingly it will be the whole job.

Naveen Agarwal, CEO NavDots; Board/C-Suite AI Advisor—Insurance, Healthcare and Financial Services

The Evolution of Insurance Asset Portfolio Optimization

ASHISH DOSHI

How machine learning is transforming investment strategies for insurance asset managers

The investment landscape is becoming ever more unpredictable driven by economic uncertainty, geopolitical risks and evolving regulations. This is putting a strain on traditional asset portfolio optimization techniques.

Traditional methods are becoming less effective in addressing the rapidly evolving financial environment and insurers are facing the challenge of struggling to balance complex regulatory and financial objectives using tools that were designed for a simpler, more stable era.

Shortcomings of Traditional Portfolio Optimization

For decades, investors have relied on techniques rooted in linear relationships such as mean-variance optimization, which seeks to balance expected return against risk. These closed-form approaches offer clear frameworks for decision-making but require simplified approximations of insurer-specific objectives.

Insurance companies face objectives far more complex than simply maximizing return for a given level of risk. They must also account for objectives such as solvency capital requirements, regulatory compliance and liquidity

Copyright © 2026 Society of Actuaries Research Institute

management. Traditional optimization approaches struggle to accommodate these objectives, particularly when constraints are non-linear and when conflicting goals must be considered simultaneously.

To overcome this challenge in managing portfolios, insurers and insurance asset managers have resorted to trial-and-error or brute-force methods, manually generating portfolios until one fits the desired criteria. While this approach can work, it is inefficient and offers no assurance of optimality. The time and resources expended in this process can be considerable and the resulting portfolios may still fall short of meeting the required objectives.

Scenario-Based Machine Learning—A New Approach

Scenario-Based Machine Learning (SBML) represents a paradigm shift in portfolio optimization, enabling users to evaluate any combination of objectives within a stochastic scenario framework. Unlike traditional methods, SBML embraces the full complexity of the real world, allowing for non-linear objectives and the simultaneous optimization of multiple competing goals.

The key to SBML is its ability to learn from vast data sets of generated balance sheet projections driven by a stochastic real-world scenario generator. Machine learning algorithms train on these projections, identifying patterns and relationships between the complex objectives and constraints. This learning process identifies asset portfolios that best meet the objectives and constraints defined in the optimization exercise creating an efficient frontier of suitable portfolios.

Targeting Balance Sheet Metrics

One of the defining features of using SBML tools for strategic asset allocation (SAA) optimization is the capacity to target the complex balance sheet metrics that matter most to insurers, leading to a targeted SAA approach.

Let's take solvency capital as an example. By and large, for all insurance regulatory frameworks globally, the amount of capital held is directly influenced by the risk profile of the investments held. Regulatory frameworks, such as Solvency II in Europe, impose strict standards on insurers, requiring them to maintain sufficient capital to cover the risks of running asset portfolios.

Under a traditional portfolio optimization approach, an efficient frontier would be generated by maximizing return and minimizing risk (in this example surplus mean and surplus cVaR), whilst searching for portfolios that meet the solvency capital requirements. This could take a while as portfolios are generated and then discarded until a portfolio is found that meets the required criteria.

SBML enables insurers to directly incorporate these considerations into the optimization process maximizing surplus and imposing a constraint on the amount of capital required. This leads to a more accurate method of finding portfolios that meet the required criteria but in a much more time efficient manner.

Copyright © 2026 Society of Actuaries Research Institute

Another example is where insurers want to consider the impact or the level of dividends payable from their surplus assets. A reliable dividend stream is important in the context of shareholders but there are several complexities that means implementing a suitable investment strategy backing this dividend stream is crucial. This is where SBML can really help, as it allows investors to specifically target the dividend as an objective metric of interest. The flexibility of the tool allows investors to either maximize or target the level of dividends whilst measuring the performance and risk profile of candidate portfolios through a stochastic lens. Adding a complex constraint such as solvency capital into the optimization can provide further insight into how to structure a suitable asset allocation.

Insurers and insurance asset managers that embrace tools that use machine learning for portfolio optimization will be best positioned to achieve their goals, adapt to new challenges, and secure their place in the evolving landscape of global finance.

Technology E&O (Tech Professional Liability)

The baseline coverage for makers. It covers financial losses suffered by customers due to the failure, negligence, or defective output of the AI technology.

As the market matures, the winners will be the companies that treat insurance not as a procurement afterthought, but as part of AI governance, contracting, and commercialization discipline.

Ashish Doshi, Head of UK Insurance at Ortec Finance

Copyright © 2026 Society of Actuaries Research Institute

Virtual Roundtable (Part 2): AI and the Actuarial Profession

EDITED BY DAVE INGRAM FSA, CERA

We convened a Virtual Roundtable to hear what people are thinking about several pressing AI-related questions. With the virtual format, we were able to get participation from actuaries in 7 countries on 5 continents! Part 1 can be found in the previous issue of the AIB. This is Part 2.

QUESTION 1. How will actuarial training evolve with the advent of AI?

- Shirley Wang (Singapore)
  What distinguishes outstanding actuaries from merely competent ones is not simply the speed at which they pass exams, but their ability to ask the right questions and recognize when a projected result does not seem reasonable.
- Yulia Nechay (Israel)
  If you can't explain why the AI's answer makes sense (or doesn't) you can't really rely on it.
- Erick Mwangi Muguchia (Kenya)
  The future actuary's value will lie not in calculating faster than AI, but in providing trustworthy judgment.
- Michael Niemerg (USA)
  I expect actuarial students will be asked to do more, see more, and operate at a higher level earlier in their careers.
- Mitch Stephenson (USA)
  Actuarial trainees will be learning how to develop, deploy, and take accountability for AI systems, even as those systems replace some of their day-to-day activities.
- Ronald Poon Affat (Brazil)
  So, the new model is a form of reverse apprenticeship.
- Mitch Stephenson (USA)
  I expect that actuarial trainees, over time, will need to 1) understand and perform tests for bias in AI models, 2) understand and implement monitoring thresholds for AI model results to check for continued reasonability, 3) understand expectations for human oversight and accountability related to AI generated output including documentation and communications, 4) be able to rely on AI assisted coding techniques and have ways to check for accuracy, and 5) understand and be transparent in documenting where AI is used in their processes, especially critical ones like financial reporting or reserving.
- Erick Mwangi Muguchia (Kenya)
  The hallmark of actuarial expertise will shift from, "I computed this," to, "I have scrutinized, challenged, and stand behind this."
- Michael Niemerg (USA)
  I also think we should expect more work along the lines of: "Check what this LLM did, compare it to what was done last year, and look for errors and changes."
- Ronald Poon Affat (Brazil)
  AI may change actuarial training because trainees may increasingly use tools and methods that their managers do not fully know yet.
- Yulia Nechay (Israel)
  AI may change how they work, but it doesn't remove the need to understand what is happening underneath.

Copyright © 2026 Society of Actuaries Research Institute

- Mitch Stephenson (USA)
  The same way we trained actuaries to use Excel ~30 years ago. Put the tools in the hands of practitioners, develop a community of practices, and share best practices through webinars, articles, and SOA-endorsed learning modules.
- Shirley Wang (Singapore)
  Encouraging actuaries to explore different tools, test alternative approaches, and ask better questions may become at least as important as teaching the mechanics behind established models and methodologies.
- Yulia Nechay (Israel)
  Even if future actuaries spend less time on the technical "black work" that trainees often do today, they will still need to understand it well enough to challenge, validate, and interpret AI's output.
- Erick Mwangi Muguchia (Kenya)
  Rather than focusing primarily on producing calculations, trainees should learn to critically evaluate and challenge AI-generated outputs using explainability tools and professional judgment.
- Shirley Wang (Singapore)
  Those who invest time in mastering AI tools, understanding their limitations, and using them thoughtfully will have an advantage in developing strong business instincts more quickly.
- Michael Niemerg (USA)
  As a result, actuaries may need basic skills in managing and updating AI knowledge bases, as well as integrating those knowledge bases into practical workflows.
- Ronald Poon Affat (Brazil)
  The master may still understand the craft, but the apprentice may understand the newest instrument.

QUESTION 2: What do you consider the biggest challenge associated with AI adoption in the actuarial field?

- Dakota Cooley (USA)
  Overconfidence in outputs that look authoritative. We're trained to be skeptical of numbers—but polished, well-formatted AI output can short-circuit that instinct. The risk isn't the model being wrong, it's us skipping the validation step because the answer looks right.
- Jason Reed (USA)
  Integration of internal corporate data assets with AI search and compute capability
- Daniel Ramsay (Scotland)
  AI governance, and specifically data governance, is the critical bottleneck. Actuaries routinely construct models using highly sensitive data: health records, addresses, protected characteristics used as risk factors. When organizational policies prohibit this data from being processed by AI systems, it severely throttles the ability of the AI to construct useful models. The technology is ready; the governance and data infrastructure often is not. Until firms resolve the tension between data sensitivity requirements and AI capability, adoption will remain patchy regardless of how capable the models become.
- Dale Hall (USA)
  Over-reliance; not enough governance and risk management.

QUESTION 3: Should actuarial professionalism standards be tightened or loosened for work with AI?

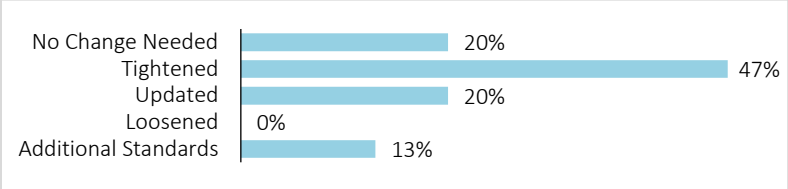

Copyright © 2026 Society of Actuaries Research Institute

QUESTION 4: How can actuaries balance the benefits of AI adoption with potential risks and challenges?

- Kathleen Locklear (USA)
  Through critical thinking and an understanding that humans must always provide needed oversight.

- Raymond Z Lai (Malaysia)
  Actuaries need to understand that AI is only a tool and we as actuaries are responsible for our outputs. Actuaries as risk manager should apply governance and control cycle in our works.

- Stephen Mildenhall (UK)
  It's still your signature on the line. Read what it writes. Understand what it does. But equally understand Luddites will surely soon be unemployed.

- Daniel Ramsay (Scotland)
  The balance is not a static position to find but a discipline to practice. Firms should run AI in parallel with existing processes before granting it autonomy, building an evidence base for trust rather than relying on vendor promises. Governance frameworks must be proportionate: heavy oversight for high-impact decisions, lighter touch for routine analysis. The profession's existing culture of peer review and sign-off already provides a natural scaffold. What matters is extending that scaffold to cover AI-generated work products with the same rigor we would apply to any junior analyst's output.

QUESTION 5: How concerned are you about "black box" issues (lack of model interpretability) in AI models?

- Simon Hatzesberger (Germany)
  One should certainly be aware of the black-box nature of certain AI models, particularly in machine learning. At the same time, it is important to understand explainable AI methods that shed light on model behavior, while remaining mindful of both their capabilities and their limitations.

- Green Chen (Canada)
  As long as it is not the core calculation, I'm not that concerned.

- Stephen Mildenhall (UK)
  Show me a person who is not a black box. And in a more evil way than any (current) model. They give misleading answers because they are trying to sell you something or actively deceive you for some reason. You rarely know why humans act as they do.

Participants are all contributors to the AI Bulletin.

Copyright © 2026 Society of Actuaries Research Institute

"Untangling" Our Wicked Problems

KATHLEEN LOCKLEAR, D. MGT., MBA

If you encountered a "wicked" problem, how would you know? And, why would you want to know? A useful way to begin thinking about this is to distinguish **wicked problems** from **tame problems.**[13]

Tame Problems

Many problems we encounter are "tame." Some defining attributes of tame problems include:

- A clear definition and consensus regarding what the problem is.
- A definitive end point (solution) that can be objectively evaluated as either "right" or "wrong."
- Belonging to a class of problems that can be solved in a similar manner.

To illustrate, we can look at this equation: $x^2 - 4 = 0$. The steps needed to solve the equation, as well as the boundaries of the problem, are clearly defined. Regardless of how many times the calculation is performed, there is only one possible (correct) answer. The computation process has a recognizable end point that is reached once the equation is solved.

Enter the Wicked Problem

Wicked problems (WPs), however, are fundamentally different. They emerge within complex adaptive systems that contain feedback loops, conflicting stakeholder interests, incomplete information, and shifting conditions. WPs have no definite stopping point because every attempt to solve them changes both the system and the problem. Climate change, catastrophic wildfire and the use of artificial intelligence are examples of WPs. Attributes of WPs are as follows:

- The problem and its broader context are continually shifting.
- Decisions about WPs must be made with imperfect information.
- There is a lack of consensus about how to define the problem.
- Intervention aimed at addressing the problem may create unintended consequences elsewhere.

These attributes make WPs particularly challenging and resistant to "solution" with traditional predictive, quantitative approaches. Success is not achieved by finding "solutions" but instead involves getting things "less wrong over time." [14]

When Historical Data Stops "Behaving"

For decades, the actuarial profession has delivered value by quantifying uncertainty and bringing structure to risk. Traditional modeling and analytic techniques have provided assurance that, by carefully analyzing historical experience, future outcomes can be mapped with reasonable confidence. A reasonable assumption has been that that the future, although uncertain, remained sufficiently connected to the past to be modeled with confidence. WPs challenge this baseline assumption. We have entered a "new normal" where conditions are unstable and cause-effect relationships are often non-linear. This means historical data has, at least for some applications, become a less reliable guide to the future. It has stopped "behaving." Paradoxically, this shift is happening at a time when our use of AI is making us increasingly dependent on data. We can begin to explore this tension by asking, "Can AI help us navigate the wicked problems that traditional models were not designed to solve?"

---

[13] This distinction was first proposed over fifty years ago by design theorists Horst Rittel and Melvin Webber. Horst W. J. Rittel and Melvin M. Webber, "Dilemmas in a General Theory of Planning," Policy Sciences 4, no. 2 (1973): 155–69, https://doi.org/10.1007/BF01405730

[14] Paul Grobstein, "From Complexity to Emergence and Beyond: Towards Empirical Non-Foundationalism as a Guide for Inquiry," Soundings: An Interdisciplinary Journal 90, nos. 1–2 (2007): 301–23.

Copyright © 2026 Society of Actuaries Research Institute

Using AI to "Untangle" Wicked Problems

Although WPs cannot be solved, we can work to "untangle" them.[15] AI is particularly well-suited for this and provides new ways to explore, simulate, and interpret WPs. Three applications of AI are especially noteworthy.

1. Pattern Recognition, Weak Signal Detection
   We are operating in an environment that contains vast and continually growing amounts of data. Using the capabilities of AI, it is possible to develop new insights about WPs. One especially promising application involves using AI to scan big data for patterns and weak signal detection. This is useful for detecting weak signals at a point when they would not otherwise be discoverable. In practice, these techniques are already being used for fraud detection, claim triage, and operational risk monitoring.

2. Digital Twinning
   "Digital twinning" involves the construction of virtual, dynamic replicas of physical objects, processes and systems. AI is already being used for this purpose, and for updating the digital twin in real time. This provides an enhanced approach for running "what if" simulations. These simulations can be used to explore cascading interactions that are otherwise difficult to capture with static models. Another benefit of digital twinning is that future failure modes can be explored without having to jeopardize "real world" assets. These capabilities are highly useful for untangling the multiple future endpoints of WPs.

3. Enhanced Scenario Analysis
   Scenario analysis has existed for many decades. However, AI can be leveraged to expand the ways that scenarios are developed and applied in practice. Using the speed and scale of AI, it is possible to generate multiple scenarios concurrently. That scenario output can be used to develop an adaptive framework for decision-making, tailored to the multiple possible futures of WPs.

Cautionary Words

It is exciting to consider the applications of AI for untangling WPs, but some cautionary words are essential, especially because we are just beginning to understand the risks embedded in AI systems and their usage.

- AI is powerful, but it is not inherently wise. Humans must provide needed judgment, institutional knowledge, and contextual understanding. AI can identify early signals, but humans must interpret those and decide what matters.
- Generative systems are vulnerable to several types of "epistemic distortions," including hallucinations, homogenization and sycophancy. At all times, humans must be active (not passive) consumers of AI output, mindful that AI sometimes produces convincing but inaccurate information.
- AI optimizes for likelihood. While this can be useful for some applications, it may also lead to the premature discounting of unlikely or improbable outcomes. Lessons learned from Black Swan Events[16] continue to have important implications as we move forward with our use of AI.
- It is important to avoid prematurely "taming" a wicked problem. This can happen when complexity is reduced early on to make the WP more "digestible." The result is a deceptive appearance of clarity that obscures deeper systemic vulnerabilities.

Concluding Thoughts

The value delivered by the actuarial profession has never been about technical precision alone. That value is much broader and involves delivering professional judgment and stewardship in the face of uncertainty. These

[15] Readers interested in an in-depth discussion of this topic can consult: Kathleen Locklear, "Wicked Problems: A Novel Approach Using Artificial Intelligence and Scenarios," Journal of Leadership & Organizational Studies 32, no. 3 (2025): 230–48, https://doi.org/10.1177/15480518251330728
[16] Nassim Nicholas Taleb, The Black Swan: The Impact of the Highly Improbable (New York: Random House, 2007), https://www.penguinrandomhouse.com/books/176226/the-black-swan-second-edition-by-nassim-nicholas-taleb/.

Copyright © 2026 Society of Actuaries Research Institute

responsibilities are becoming even more important in AI-assisted environments. Increasingly, the future value of actuarial science may lie in supporting organizations as they seek to remain adaptive within complex, continuously changing systems and in the face of WPs. Doing this effectively will require combining quantitative rigor with other tools including systems thinking, scenario planning, ethical oversight, and interdisciplinary collaboration.

Dr. Kathleen Locklear is Associate Professor of Risk Management and Insurance at the State University of New York at Oswego.

Editorial Committee

Jon Forster, ASA
Dave Ingram, FSA
Jing Kai Ong, FSA
Ronald Poon Affat, FSA
Frank Quan, PhD
Lisa Schilling, FSA
Darryl Wagner, FSA

Thank you for reading the September 2026 SOA Research Institute AI Bulletin. We hope you found these insights valuable. Stay tuned for future editions as we continue to explore the evolving landscape of AI and its impact on the actuarial profession. We encourage you to engage with the SOA Research Institute and share your own experiences and perspectives on AI. For questions, comments, and article submissions, contact Research-AIB@soa.org.

Copyright © 2026 Society of Actuaries Research Institute

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, Suite 1000
Chicago, IL 60631
www.SOA.org

Copyright © 2026 Society of Actuaries Research Institute