

ACTUARIAL INTELLIGENCE BULLETIN



July 2025

TABLE OF CONTENTS

The AI Enhanced Actuary	2
Agentic Artificial Intelligence.....	3
From the Frontline: Big Data and AI in Healthcare.....	5
AI Uncovers M&A Appetite in the Filings of Private Insurers	7
Shipwrecks and Spreadsheets: Steering AI with a Human Compass	9
Actuaries No Longer Have an Excuse: Help Writing Articles with AI.....	11
The Risk Framework That’s Changing the AI Conversation.....	13
Enhancing Data-Cleaning with Generative and Predictive AI	15
AI in the Determination of the Actuarial Discount Rate	17
AI Insights Meeting May 2025.....	19
Actuarial AI Links	21
Are You a Typist or an Actuary?	22
Editorial Committee	23
About the Society of Actuaries Research Institute	24

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.

The AI Enhanced Actuary

Excerpt from Opening Address of 2024 SOA AI Insights Meeting

RONALD RICHMAN

I think society benefits by having AI enhanced actuaries, actuaries who can build models more efficiently, provide better services through higher accuracy and higher speed, we can contribute to more stable and resilient insurance markets, also from the perspective of personalizing pricing, understanding what is an actually fair premium, and also ensuring that we're not inadvertently discriminating against policy holders. I think those are some obvious policy holder advantages, together with some of the advantages you get from applying Gen AI systems like hyper

“Something that's emphasized again and again is understanding your models, and that understanding not only relates to model methodology, but understanding exactly how decisions are being made. Applying some of the approaches like post hoc interpretability or using inherently interpretable models, gets you quite far along those requirements, and then model validation and understanding the data going into these models is another key issue.”

personalization, which can radically improve customer experience. There's an obvious value, proper value proposition for insurers, where insurers can gain from increased efficiency innovation and better data driven decision making. And I think most importantly for this audience, there's also a great opportunity for us to enhance our skills, lead in the AR space and drive innovation. And if we get this right, I think it's not too far-fetched to think that actuaries could come to be seen as experts even in the AI domain. So in other words, outside of our traditional actuarial space, where we have a great understanding of risk, we are committed to professional and ethical standards, and with newly acquired AI skills and hopefully deep technical skills, then we could perhaps play a pivotal role in shaping what the implementation of AI looks like in wider society, not just in our traditional domains, but we first need to get right some of the challenges about how we do this in our traditional domains. Perhaps you'd have huge actuarial roles,

like an AI risk assessment specialist or resilient AR systems architect. After getting this right in a highly regulated and complex domain like insurance or asset management or banking, it would very easily then allow actuaries to generalize that experience and create these sorts of future roles to do this, and this is for those of us who are involved within our actuarial organizations like the SOA or the IFO, a or ASA, we need to think about, how are we educating our next generation of Actuaries in a lot of jurisdictions? The current state of the actuarial syllabus is that these AI skills are under review with a proposal to incorporate the concepts more thoroughly. But the problem is that this is still transformation in progress while it feels like other areas and other industries are running ahead of us. So, I think a lot needs to be done to get this ready sooner. And if anyone is interested in an example syllabus for actuaries you might be teaching, or for future syllabus reorganizations. About what that could include. So, let's think about positive and negative scenarios. I think the positive scenario, which I hope what is what we see, is that in the coming years, we see thriving AI enhanced actuaries. This would involve a rapid adaptation to AI and ML through modernized education, more innovation in applying AI within our traditional external problem set, developing data driven thinking, together with an ethical AI implementation, making sure that we stick to our professional guidance. And this would allow actuaries to lead industry innovation and not almost fall by the wayside with other professionals. I think the negative scenario is if we fall behind and through resistance to embracing AI and deep learning tools, we really face the risk of a loss of our technical domains to other professionals like data scientists or machine learning engineers, and we could find that the profession is confined to a more limited space, maybe only looking at regulatory matters. I think this would be a great pity, given the depth and potential of the actuarial skill set in an AI driven world. So, the answer, in my mind, is this vision of the AI enhanced actuary. And how could this

be done by having AI enhanced actuaries who use AI powered models to provide more accurate insights, faster modeling, better data driven analysis. And this would allow us to amplify the unique strengths we already have together with these newly emerging tools. As we've said, there are challenges along the way, model explainability, regulatory compliance and some risks around bias and discrimination, but if one looks around in the technical actuarial literature, a lot of solutions already exist. I think that accomplishing this would lead not only to value creation in our traditional markets but also open up the opportunity for actuaries to shape AI implementation in wider society. And if we do this, I think the future of the actuarial profession is really bright, and it will start with how we start to make these choices going forward.

Ronald Richman, FASSA CEO insurAI



Agentic Artificial Intelligence

NORMAN STORWICK

At its simplest, an AI agent is a program that uses AI to perform a specific task or solve a particular problem. It likely has an appropriate skill set and can use the tools and information available to it to accomplish its goal.

Modern AI agents extend these capabilities and have two primary attributes:

1. They can call a large language model (LLM) to plan, act, and complete some tasks. It may start “You are an expert underwriter.” or “You are an actuary with years of experience calculating manual rates for large group health insurance.” The content of this prompt is often vital to the success of the agent. Thorough descriptions of the agent’s role, suggestions of project plans, and even examples of formatted output can improve the agent’s performance.
2. They may have tools defined which the LLM can call. These tools give the LLM capabilities it doesn’t have on its own. These can be specific calculations which the LLM cannot easily replicate, access to unique data, or even the ability to interact with the real world. Any action you can program can be adapted as a tool available to an agent.

For our underwriting example, we built two tools for our agents. One performs manual rating and the second uses credibility weights to blend manual and experience rates. By using tools, we can ensure the agent will follow our company’s processes and that the results will be 100% consistent.

We identified three specific roles in the underwriting process and created agents for those roles. The roles were:

1. Underwriter—This is the project manager. It understands the roles of the other agents and can determine when to call them.
2. Medical underwriter—This agent has deep understanding of healthcare and disease processes. It can review the diagnoses of high-cost members and predict which will be high-cost next year.
3. Actuary—This agent knows how to calculate manual rates and apply the credibility factors.

Of course, these agents are patterned after roles in underwriting teams.

The last element of a multiagent AI solution is an interaction framework. These frameworks support conversations among the agents that range from free flowing “round robin” conversations to tightly managed conversations where Agent A needs to complete a step and then pass the spotlight to Agent B. For our example, we used a round-robin model where all the agents were sitting around a virtual table following the guidance of the underwriter agent through the steps in the process.

To test our framework, we presented the agents with a description of the prospect, a census file (age and gender of all the beneficiaries), a high-cost claimant report, and simulated claims experience. The underwriter then initiated the process.

The underwriter agent deduced the prospect’s industry, then provided the census information, census, industry code, and trend information to the actuary agent. The actuary agent applied the manual rating.

With a manual rate, the underwriter agent analyzed the claims data and calculated an experience rate. The medical underwriter agent reviewed the high-cost claims report. It recognized which conditions were unlikely to recur and which were likely to continue. It provided estimates of the future costs of high-cost claimants. The underwriter adjusted the manual rate and then calculated a credibility-adjusted average of the manual and experience rate.

“Next, we plan to fine tune a foundational model to provide better predictions about high-cost claimants. And fine-tune another AI model to perform executive approval of the quote to ensure it is consistent with the company’s underwriting strategy.”

Amazingly, this whole process ran in a matter of seconds. We reviewed the conversation between the agents. These conversations were conducted in English and gave us the chance to be a fly on the wall of the agent’s virtual office. We read where the agents ask for information or complete their step.

As we developed the application, reviewing these conversations gave a great view into its workings. Much like managing a team (and very much unlike debugging a computer program) evaluating communication among the team can reveal problems in the process. Here, the conversations were loose, and the team often failed to complete the task. Often, the actuary agent would say that they wanted the census file to calculate the manual rate. Before another agent could provide them with the census file, the actuary agent asked again. And again.

Empowering the Underwriter agent to manage the process solved this problem. The Actuary learned to wait for the Underwriter to provide them with everything they need. With these changes, the process improved, and the team was able to successfully, and accurately, calculate the quote.

What we find most enticing about this approach is that it can improved further. For instance, all the agents use GPT 3.5 to make judgements and guide their actions. But GPT 3.5 wasn’t developed to be an expert in actuarial processes or medical underwriting.

AI vendors offer tools that allow you to fine tune these foundational models to better solve your problems. You can provide a data set to the foundational model and adjust the model’s parameters, so it is better able to respond to your problems. Next, we plan to fine tune a foundational model to provide better predictions about high-cost claimants. And fine-tune another AI model to perform executive approval of the quote to ensure it is consistent with the company’s underwriting strategy.

Looking ahead, we see tantalizing opportunities. A robust underwriting solution would include additional agents that provide peer review, complete claims analysis to identify trends, recommend care management programs, predict

their impact, and structure the quote to match the preference of the broker. Each of these agents can be using an AI model that has been uniquely trained for their purpose.

Norman Storwick, FSA Chief Actuary and VP of Underwriting Curative



ACTUARIAL INTELLIGENCE BULLETIN

From the Frontline: Big Data and AI in Healthcare

IAN DUNCAN

I recently chaired two panels on big data and AI at the Arbital Health Value-based Healthcare Summit. What follows is a summary of the key themes and insights that emerged—presented in the words of the panelists themselves.

Panel participants represented a broad spectrum: clinicians, the head of benefits at a large employer, the founder of a provider group for the underserved, and representatives of a provider enablement platform company, a behavioral health provider, and cancer management and medication therapy optimization companies.

“We are looking for small data—what are those few parameters that actually give a signal?”

Do We Have the Right Data?

I opened by asking whether we have too much data—and whether we have the right data.

Kevin Holub, VP of Strategy, Counterpart Health

“You ask whether we are inundated, and I say, no, I still want more data. And I’d like easier access to data in more normalized and standardized ways. I love FHIR and I love EHRs that say they are compliant with FHIR, but most of them, when you go under the hood, they’re not. We believe that there is no limit to the amount of data that is helpful. The key is turning all data into something that is helpful. CMS is doing some great things on interoperability and the BCDA is wonderful, but basically no one else uses it other than CMS.”

Signal vs. Noise

The panel agreed that while they are flooded with data, the real challenge is separating signal from noise.

Jon Bloom, MD, Podimetrics

“We are looking for small data—what are those few parameters that actually give a signal?”

The Last Mile in Healthcare

The panel agreed, however, that data and AI are not sufficient. Insights need to be delivered to patients to provide care and change behavior.

Jim Wallace, DBA, CEO, Decision Rx

“How do patient-provider interactions affect the outcomes? The conclusion is that 37% of variation in outcomes is directly attributable to whether the doctor believes the patient and the patient believes the doctor.”

“Wall Street thinks AI should be at the back end at the last mile, just using a large language model to concisely compress what a doctor should say to the patient. If you don’t have the right emotive cognitive layer, it doesn’t matter what you say, the patient’s not going to hear it. We use big data on the front end to figure out what’s predictable and what we should use with the provider team who are acting on the data to steer to the right conclusions.”

Terri Maxwell, PhD, RN, Palliative Care Specialist

“Ten years ago, we developed our model to identify people at risk for over-medicalized death, to engage them in a program and then deploy palliative services to them. We didn’t have a lot of clinical information, however. We overlaid the clinical data manually on the results of the predictive model. But immediate access to clinical data would have been a game-changer.”

The panel would like more data, more timely data and a way to make the data more insightful.

Using AI in Practice

I then asked how they are using AI in practice.

Terri Maxwell, PhD, RN

“I’m advising a company that developed a palliative care app that’s very AI-driven, and I have evolved my thinking on this. I’ve really become bullish on the degree to which AI can successfully engage with people in ways that I think sometimes are better than even very experienced clinicians. In terms of goals of care conversations, normally you would think you need to be trained. Well, I’ve seen examples where the AI bot is picking up cues to help time the goals of care conversations.”

Jon Bloom, MD

“It’s funny, I watched my dad talk to Siri. He’s struggling. He’s not making it work very well. In my company, we typically care for seniors and there’s a lot of skepticism. AI isn’t going to take over, but it is catching up. And as I have a conversation, it’s suggesting empathy back to me. It’s having a very rich conversation with me. And I know it’s a piece of technology, it’s ChatGPT. But it’s crazy how fast it’s getting there.”

Will AI Replace the PCP?

At the end of the panel, I asked the audience: *Will AI eventually replace primary care physicians?*

Many core physician tasks—diagnosis, prescribing, ordering and evaluating tests and treatments, and referring to specialists—are governed by clinical guidelines. If those guidelines can be integrated with AI and conversational tools, an algorithm could potentially take on much of the PCP role.

As several speakers noted during the panel, ChatGPT is making impressive strides in its ability to hold a conversation, pick up cues, and suggest empathy. It may be difficult for the public to accept technology in place of the PCP, but with a projected shortfall of 187,000 primary care physicians by 2037, this may become inevitable.

I conducted an informal audience poll later—the answer was 3 to 1 in favor of AI replacing the primary care physician.

Since the panel, I have discovered two companies developing applications in this space. If any readers know of other companies, let me know!

Ian Duncan, PhD, FSA, FIA is Senior Advisor at Arbutus Health and Adjunct Professor at UC Santa Barbara



ACTUARIAL INTELLIGENCE BULLETIN

AI Uncovers M&A Appetite in the Filings of Private Insurers

TOM MASON

The Management Discussion and Analysis sections of annual statutory statements filings offer a wealth of information on private insurance companies—in this case, a source of potential buyers and sellers. But such information has historically been hard to obtain for a few reasons. One is that the documents are not publicly available and must be purchased from the National Association of Insurance Commissioners. Another is that the documents come in PDF, which requires sophisticated data science techniques to make them machine-readable. Finally, there is a cost/benefit analysis. There are thousands of documents that need to be parsed, a time-consuming exercise where the hit rate is very low: Only 3% of the 3,145 documents contained the information we were seeking.

But with modern technology, we automated these processes in Python using advanced PDF extraction tools and a large language model. One of the advantages of using LLMs for natural language processing is that they understand related terms. This was especially helpful for life insurers, since the LLM identified many instances where the company was interested in reinsuring blocks of business, even if the exact words “merger” or “acquisition” were never written in the document. These block deals are commonly considered akin to acquisitions in the life insurance industry and can be very large in size, sometimes rivaling the price tags of whole-company acquisitions.

About The Analysis

We created a pipeline that extracts the text from a PDF document and feeds it to a large language model that can read it and answer questions. To build the pipeline, we started with a folder that contained all the raw PDF files and looped through it, extracting the text from each file and having the LLM answer the question based on just that document. S&P Global built an in-house tool, Spark Assist, which leverages OpenAI's GPT-4o mini model and has built-in functions for uploading documents programmatically and asking questions of that document. Spark Assist is also accessible via an application programming interface, allowing us to automate the process of reading thousands of documents.

“Based on our time trials, it would take approximately 150 hours for a human to read 3,145 documents and find information in them, whereas our Python pipeline can do it in about two hours.”

Based on our time trials, it would take approximately 150 hours for a human to read 3,145 documents and find information in them, whereas our Python pipeline can do it in about two hours. What's more, we can ask the LLM as many questions as we want. Our Python script accesses the LLM via an application programming interface, so whenever we have a new question, we simply enter the question, hit run, and the program will read all the documents again and answer the new question we've given it for each document. For instance, we have asked other questions like, “Does this text mention any forward-looking commentary?”, “Does the company mention reinsurance renewals in this text? If yes, what does it say?” and “How did the company change its investment strategy in 2024?”

A Spectrum of Intentions—What Insurers Are Saying

Several insurers expressed a specific desire to engage in acquisitions. One P&C firm focused on commercial lines, for example, stated it is “actively open to another acquisition opportunity, with particular emphasis placed on a commercial lines affiliation and/or potential niche commercial lines book of business.” Another aims to “continue to seek opportunities to acquire existing books of commercial auto business.” In the life insurance sector, acquiring blocks of business through reinsurance is a common growth strategy. Some companies explicitly see M&A as a path to building scale and diversifying their operations.

The analysis also identified companies on the other side of the M&A equation. Commentary revealed that for some, a “receiver is negotiating the sale of the company to a potential buyer,” while another's “ownership has decided to sell the company and is in negotiations to finalize the transaction.” Some smaller entities are “seeking additional investors and/or companies interested in acquiring the company.” Restructuring plans also surfaced, with one P&C company's future options including “a merger into one of the other AXIS US insurance companies, a sale of the shell,” and another's management planning to eliminate a company from its corporate structure through merger or sale.

The underlying motivations expressed for this M&A interest are diverse. Growth in specific business lines, such as medical professional liability, is a driver for some. Others are eyeing market expansion, including entering new geographic areas or acquiring technical service companies to access new markets. For some mutual insurers, M&A is seen to improve financial metrics like the spread of risk and expense ratios. Even fraternal benefit societies are looking to M&A, with one stating its strategic plan includes remaining “open to merger possibilities and to continue forming strong alliances.”

In conclusion, the application of AI to parse these extensive private filings has peeled back a layer of obscurity, revealing a dynamic M&A landscape. The varied commentary, from active pursuit to cautious openness and strategic sales, underscores ongoing repositioning within the U.S. insurance industry. This AI-driven approach not

only demonstrates remarkable efficiency but also provides invaluable market intelligence crucial for actuaries and other stakeholders navigating this complex sector.

An earlier version of this article appeared in S&P Capital IQ on 20 May 2025.

Tom Mason Principal Quantitative Research Analyst S&P Global



Shipwrecks and Spreadsheets: Steering AI with a Human Compass

CAROLINE MADRAK

Artificial intelligence poses an essential—and uncomfortable—question: If an Excel sheet can now learn to think, what will be the role of tomorrow’s actuary?

AI is no longer the stuff of science fiction, it’s already changing the landscape of insurance, risk modeling, and beyond. And while many of these technologies have existed in some form for years, their rapid evolution has policymakers questioning if regulation has kept pace. That’s why it’s crucial for our community to get comfortable—not just with the capabilities of AI, but with the conversations around it. Because, like every great technological leap, progress arrives with its darker twin: risk. There were no shipwrecks before ships.¹

As the regulatory environment continues to evolve, ethical principles and individual integrity must be our compass. As actuaries, we cannot just assume that because a model can be built, it should be. That kind of discernment—nuanced, self-aware, and often uncomfortable—is part of what will distinguish the human professional from the machine.

At the heart of responsible AI are a few foundational principles all actuaries should understand:

- Fairness: AI systems should not unfairly discriminate. But fairness also is not an absolute metric, what feels fair to one person might not to another. It’s inherently subjective, context-driven, and culturally informed. That does not make it impossible to aim for—it just makes it impossible to ignore.
- Safety: Systems must be robust, accurate, reliable, and secure. In other words, they should work as intended, even under stress, and protect sensitive information. No model is perfect, but the bar for safety must be high when lives, health, and livelihoods are on the line.
- Transparency and Accountability: Trust is foundational to the success of AI. We need AI systems that are explainable, auditable, and subject to real oversight. If a decision cannot be explained, should it be implemented, especially in fields like insurance where those decisions carry real weight?

These values are not new to our profession. In fact, they echo core actuarial principles embedded in the Code of Professional Conduct and our established modeling standards. The need for fairness, safety, transparency and

¹ “The invention of the ship was also the invention of the shipwreck...Every technology carries its own negativity, which is invented at the same time as technical progress.” Attributed to Paul Virilio, *Politics of the Very Worst*, interview by Philippe Petit, trans. Michael Cavaliere (New York: Semiotext(e), 1999), p. 89.

accountability aligns directly with Precepts 1 and 3, which emphasize integrity, competence, and the importance of communicating clearly to maintain public trust. In many ways, the growing presence of AI simply extends the reach of these ethical commitments into new, more complex territory—one where our professional judgment and sense of duty become even more essential.

These values need to be applied across the full actuarial modeling cycle because risk and bias can sneak in at any step:

- Data: Raw data is a myth. Data is never neutral—it is created, selected, and shaped through human decisions. The tools we use to collect it can introduce noise, distortions, or systemic gaps. Bias does not start in the model; it starts in the spreadsheet.
- Model: A model does not know how to discriminate—and it also does not know not to. If a pattern in the data improves the objective function, it will be captured and reinforced, whether it is just or unjust. That is why it is critical to ask not only, is it accurate? But also, is it right?
- Implementation: Even the most thoughtful model, perfectly trained and tested, can cause harm when deployed in the real world. Implementation is where unintended consequences bloom. Sensitivity testing and scenario planning are vital for spotting weak spots before they become news stories—or worse, lawsuits.

“Do good. Cause no harm. Give everyone their due.”

It might help to borrow some wisdom from the past as we sail forward into the unknown. Ulpian, the Roman jurist credited with the first recorded life table, once wrote: “Juris praecepta sunt haec: honeste vivere, alterum non laedere, suum cuique tribuere.” (“The basic principles of law are to live honorably, not to harm any other person, and to render each his own.”)

In other words: Do good. Cause no harm. Give everyone their due.

Those values hold—maybe even more so now, in this age of algorithmic decision-making. Because while AI may be able to calculate a thousand times faster than a human, it still does not understand what it’s calculating for. It does not know the weight of a life, or the dignity of a person. It does not feel the consequence of a false negative or an unfair premium.

That is our job. That is still the human part.

And yes, there will be shipwrecks. We are navigating waters no map has charted. But if we know that risk is inevitable, we can prepare. Those principles—fairness, safety, transparency and accountability—are not just buzzwords. They are the life rafts we should insist on bringing aboard.

In the end, AI might change what it means to be an actuary—but it will not replace the core of why we do this work. Because no matter how advanced the models become, or how efficient the systems grow, the role of the actuary is not just to quantify risk. It is to remember that behind every calculation, there is a human being.

My views on this topic have also been shaped by the SOA’s former Executive Ethical and Responsible Use of Data and Predictive Models Certificate Program.

Caroline Madrak, FSA, VP, Senior Client Manager, Swiss Re L&H Global



Actuaries No Longer Have an Excuse: Help Writing Articles with AI

MARY PAT CAMPBELL

On May 27, 2025, the Society of Actuaries hosted a webcast titled “AI Tools You Can Use,” moderated by David Schraub and featuring presentations by Louis Abraham and me, Mary Pat Campbell. The session explored the intersection of artificial intelligence and actuarial practice, emphasizing both the opportunities and responsibilities that come with adopting new technologies.

Did you like that first paragraph? It was composed “by” Microsoft CoPilot, after some prompts, using our slides from the webcast. (With one edit – can you guess my edit?)²

Here are the prompts:

Prompt: Could you make a short article based on this presentation: MPC AI Tools 2025_05_27 v6.pptx

[Response from CoPilot, which was way too short and generic]

Prompt: I need about 1000 words for the article. It will be for an actuarial newsletter to let actuaries know how to use AI tools to write articles

[Longer response, with downloadable text file, but now it’s extracted text from my slides ... a lot just filler]

I took the opening paragraph from CoPilot, so you can see what that sort of output looks like. The webcast was recorded, and you should check it out.³

You might not be interested in my section on using AI for communication tasks like articles, PowerPoint decks, and LinkedIn posts. Louis Abraham’s segment covered GenAI generating Python code to analyze datasets from scratch, regardless of familiarity with the data. He demonstrated this live with 2024 Medicaid drug utilization data.

AI Provides Communication Access – for Actuaries, too!

In 1988, my Aunt Mary Pat Radabaugh, then Director of IBM National Support Center for Persons with Disabilities said, “For most people technology makes things easier. For people with disabilities, however, technology makes things possible. In some cases, especially in the workplace, technology becomes the great equalizer and provides the person with a disability a level playing field on which to compete.” At the time, she had been working within IBM on technologies such as “talking computers” (text-to-speech), as speech transcription was in a nascent stage.

Current AI technology offers extensive access. I have chronic pain and neuropathy in my hands, which limits my typing. I use Otter, an online transcription tool, to record, transcribe, and summarize actuarial work groups. Sometimes, I book a meeting with myself to transcribe spoken articles.

A short list of the tools I regularly use:

- Otter: transcribing audio, especially regular meetings, summarizing
- Microsoft CoPilot embedded in Office
- PowerPoint Designer (with pre-built Templates and Themes)

² I added “me,” as I did not tell CoPilot that I was one of the presenters and should be referred to in that manner.

³ Available at <https://www.soa.org/prof-dev/recordings/2025/april/ai-tools/>.

- Met Museum Digital Collection⁴

The last one doesn't have AI (yet), but it is a truly vast art and historical collection, most of which has public domain images. GenAI images are not as compelling as human-created art, especially the art particularly curated by the Met, the best museum in the world.⁵

Use the tool specific to your own need

In my own case, my problem is not writer's block or even the difficulty in writing to begin with.

It's overwriting.

I dug into my collection of drafts of articles for SOA newsletters in the past—I easily went over editors' stated limits. I was given a 1,000-word limit for this piece. So, you can see how I use AI normally, I wrote a normal draft for me. My first draft clocked in at 1,227 words.

That would never do.

With Microsoft CoPilot built into Word, I can select text that is particularly flabby and ask it to rewrite. This is my most common use of MS CoPilot.

Original

If you are publishing results with your name on them, you need to review them. This does not mean there will be zero errors; that was not the case when one had non-AI ways of production. However, we know of certain ways these LLMs and similar systems go awry... it would do well that we not give up our reputations. This is no different from when actuarial systems moved away from desktop calculators and mainframes to the era of spreadsheets and laptops.

Rewrite with CoPilot

If you are publishing results under your name, review them. Errors are inevitable, even without AI. We know these systems can fail, so protect your reputation. This is like when actuarial work shifted from calculators to spreadsheets and laptops.

I consider this part of my toolbox of writing tools, alongside spellcheck and grammar checking tools.

After using CoPilot in Word, I was able to slim down my draft to 988 words (according to Word's official count.)⁶

While ChatGPT, Gemini, Grok, and similar generic AI tools are interesting, I find the special use AI tools, such as CoPilot embedded directly within Word or PowerPoint, to be more useful in helping me in terms of productivity, for day job, actuarial volunteering, and blogging.

It's Just a Tool – You're Still Responsible

Here is the Society of Actuaries' policy on AI-generated content in authored pieces for section newsletters:

⁴ <https://www.metmuseum.org/art/collection> Provide link to webcast page here, or however SOA does it now.

⁵ Yes, I'm biased.

⁶ The final published article is likely to be different.

AI Generated Content

Authors are required to take responsibility for all AI generated content. AI-based tools and technologies, including but not limited to large language models (LMs), generative AI, and chatbots (i.e., ChatGPT) cannot be listed as an author. Authors are required to acknowledge where AI tools are used, document the AI generated content accordingly, and continue to sign Society of Actuaries publishing agreements.⁷

I take full responsibility for this article's content. Although CoPilot generated some text, I reviewed and edited everything to ensure appropriateness. The initial file given to CoPilot was mostly compiled by me.

When I used CoPilot to edit down some of my content, CoPilot took my original text as its starting point and gave me three options to choose from. I made my choice. Then I edited it further.

It's a process.

There have been instances where AI-generated content, such as legal briefs, included incorrect or fabricated court case references. It is important for users to review the output thoroughly to ensure accuracy. This is comparable to running models without reviewing results and implementing controls.

Review what you will publish. Errors are inevitable, even without AI. We know these systems can fail, so protect your reputation.

Why AI? A Force Multiplier

In an era of noise, it is difficult to be heard and to influence. Actuaries may be drowned out by those who are more facile in their communications.

AI can help us have a fighting chance to get our message out. In my case, it can help my edit down to the core message, and then re-purpose that message for multiple platforms.

You may have a more modest goal in mind, perhaps an internal audience in your company. Or, hey... maybe you could write an article for this very newsletter! Why not try it out?

Mary Pat Campbell, FSA, MAAA Vice President Insurance Research, Conning



**ACTUARIAL
INTELLIGENCE
BULLETIN**

The Risk Framework That's Changing the AI Conversation

*An Interview with Ronald Poon Affat, FSA, FIA, MAAA, author of **AI Risk Management Frameworks: An Expert Panel Discussion***

Editor: Ronald, your new report *AI Risk Management Frameworks: An Expert Panel Discussion*, published by the Society of Actuaries Research Institute, brings together perspectives from regulators, academics, and industry leaders. What motivated you to take on this project?

⁷ Source: <https://www.soa.org/publications/writing-ops/section-author-guidelines/>, retrieved 28 May 2025.

Ronald: What really drove this project was the growing realization that AI is no longer something “on the horizon.” It’s already shaping underwriting, pricing, customer service, and capital management. But very few actuaries feel equipped to evaluate AI systems, let alone challenge their risks. I wanted to shine a light on how professionals across sectors are learning to manage this—what frameworks they use, what blind spots they’ve encountered, and what advice they have for others in our field. A key focus of the discussion was NIST’s AI Risk Management Framework.

Editor: For those unfamiliar, what exactly is NIST and why are they involved in AI?

Ronald: NIST stands for the **National Institute of Standards and Technology**. It’s a U.S. government agency under the Department of Commerce, and for decades it’s been the go-to body for setting measurement standards across science, cybersecurity, and now AI.

Editor: So, what does this NIST framework actually do—and why should actuaries care?

Ronald: Their **AI Risk Management Framework (AI RMF)** was created to help organizations develop and use AI responsibly, even if they’re not technology companies. It’s voluntary but very practical, and increasingly influential. In fact, **Microsoft** has publicly stated that it’s aligning its AI governance strategy with NIST’s framework, which speaks volumes about its credibility.

Think of it as the actuarial control cycle for AI. The NIST AI RMF breaks AI risk down into core functions—**Govern, Map, Measure, and Manage**—with detailed guidance under each. It asks organizations to evaluate issues like data quality, model transparency, and unintended harms. You don’t need to be a data scientist to use it. You just need to ask the right questions: What is this model supposed to do? Who might be harmed by it? Is it working as expected? That’s exactly the kind of risk thinking actuaries are trained to do. The framework gives us a common language to engage with data scientists, compliance teams, and executives.

Editor: Let’s talk about bias. You make it a central theme of the report. Why?

Ronald: Because bias is baked into the data. AI models are trained on historical data, and that data includes **both helpful and harmful biases**. A helpful bias might be charging higher premiums for smokers—there’s evidence behind it. But harmful bias might emerge if ZIP codes are used as proxies for race or income, creating systemic unfairness. The AI doesn’t understand justice—it just learns from patterns. That’s where actuaries come in. We need to recognize when those patterns are replicating inequality and take **corrective action** to protect clients and ensure fairness. The most important takeaway is that **AI is not neutral**. If we treat it like it is, we risk automating past discrimination.

Editor: Section 5 of your report shares reflections from professionals who are deeply involved in AI governance. What stood out to you?

Ronald: Their honesty. These are highly experienced people—some leading AI risk at major institutions—and yet every one of them admitted how humbling the learning curve has been. They talked about overestimating the technical fixes, underestimating the human factors, and learning to navigate ambiguity. What really stuck with me was the idea that **AI safety is not just about engineering—it’s about collaboration, ethics, and context**. That’s something actuaries are well positioned to understand.

Editor: Some actuaries might say, “AI safety isn’t my lane.” How would you respond?

Ronald: I'd say: It must be your lane. Actuaries are already expected to validate assumptions, challenge model logic, and think probabilistically about uncertainty. That's exactly what AI safety needs. Add in our professional standards and public interest responsibilities, and actuaries can serve as the **second pair of eyes**—not just for financial solvency, but for fairness and long-term societal impact. The danger isn't just in rogue AI developers. It's in well-meaning professionals who assume someone else is taking care of the risks.

Editor: For actuaries who feel behind on this topic, where's a good place to begin?

Ronald: Start small. Download the **executive summary** of the NIST AI RMF—it's short, readable, and actionable. Ask if your company has an AI or model governance committee. Volunteer. Follow credible voices in the field—academic researchers, government publications, insurance regulators. You don't have to become a programmer. But you do need to develop enough fluency to **ask the right questions—and recognize when the answers fall short**.

Editor : What's your main hope for readers of this report?

Ronald: That they'll see AI safety not as a future problem or a tech team issue, but as a live professional responsibility. The report doesn't pretend to have all the answers. But it shows how real professionals—across the insurance industry, regulation, and academia—are starting to engage with this challenge. I hope actuaries see themselves in those stories and realize you don't need permission to get started. The industry needs your perspective—now more than ever.

About the Expert Panel

Elaine Newton, PhD Principal for AI Standards & Tech Policy at Amazon Web Services (AWS).

Dorothy Andrews, PhD, ASA, MAAA, CSPA Senior Behavioral Data Scientist & Actuary at the National Association of Insurance Commissioners (NAIC).

Dave Ingram, CERA, FRM, PRM, FSA, MAAA Risk Management Actuary and Society of Actuaries Board Member.

Shane Leib, FSA, MAAA Director of Actuarial Research at Moody's Analytics and Assistant Teaching Professor at the University of Notre Dame.

Ronald Poon-Affat, FSA, FIA, MAAA, CFA, HIBA (Moderator) Reinsurance specialist and Independent Board Member.

Download the full report:

<https://www.soa.org/resources/research-reports/2025/ai-risk-management-frameworks/>



ACTUARIAL INTELLIGENCE BULLETIN

Enhancing Data-Cleaning with Generative and Predictive AI

JEFF HEATON

Accurate and reliable data is critical for actuarial modeling and decision-making in insurance. Actuaries regularly face challenges with inconsistent formats, missing data points, and implausible or erroneous records. It is important to deal with these inconsistencies early in the research pipeline and cycle back as additional issues are uncovered. While traditional deterministic cleaning methods help standardize data, new developments in generative AI (GenAI) and predictive AI offer innovative ways to further enhance data quality. AI systems can come from a variety of vendors; as a result, it is critical always to protect data privacy.

Deterministic vs. Predictive Data-Cleaning

Multiple methods allow AI to help clean data, including both deterministic and predictive approaches.

Deterministic cleaning involves clear, rule-based data transformations. AI can assist with the creation of these rules, or they can be hand-coded by actuaries and data engineers. Common examples include formatting dates consistently, normalizing policyholder addresses, and standardizing units such as currency or measurements. Because these processes follow explicit, predictable rules, actuaries trust deterministic cleaning methods to deliver precise results.

Predictive data-cleaning, by contrast, leverages predictive AI to manage uncertainty. Actuaries frequently encounter datasets with missing or suspicious values that cannot be deterministically fixed. Predictive AI methods such as regression models, decision trees, or neural networks estimate these unknown values based on historical patterns. For example, predictive AI can estimate missing claim amounts or incomplete demographic information using previously observed trends. Actuaries must recognize that these solutions are probabilistic, requiring judgment to interpret their accuracy and reliability.

Predictive AI for Missing Data and Anomalies

Predictive AI is particularly adept at identifying missing or anomalous data. For instance, predictive models trained on historical claims data can accurately estimate missing fields such as age, occupation, or vehicle mileage. These AI-driven estimates typically include confidence measures, indicating the reliability of each prediction. Such insights help actuaries integrate these probabilistic predictions into modeling workflows. Often, predictive AI can become an additional set of eyes, identifying issues within the data that need to be addressed before the project can continue.

GenAI for Identifying and Filtering Unreasonable Data

GenAI, particularly large language models (LLMs), introduces advanced capabilities to detect anomalies or questionable records. GenAI analyzes complex, unstructured information, such as claims notes or underwriting comments, to identify inconsistencies, signs of fraud, or implausible scenarios. For example, a GenAI model can review claim descriptions and highlight entries that appear unreasonable or inconsistent with historical patterns.

However, actuaries must use judgment in deciding whether to discard or further investigate flagged records. GenAI can suggest removing certain rows, but actuaries must assess these recommendations carefully, considering business context and applying domain expertise.

AI Estimates vs. Hallucinations

So how can actuaries distinguish predictive estimates from GenAI hallucinations? Predictive estimates rely on patterns derived from historical data, and actuaries can evaluate their reliability through statistical validation. For example, when predictive AI imputes a missing claim amount, it bases the estimate on prior observed values, providing measurable confidence intervals.

In contrast, AI hallucinations are plausible but entirely fabricated outputs generated without grounding in actual data. These occur primarily when sufficient context is not available, leading the GenAI to produce coherent sounding, yet incorrect, information. Actuaries should always apply rigorous validation methods to GenAI outputs, ensuring that no hallucinated data points enter critical models or analyses.

When AI is in the loop, human actuaries must remain vigilant in checking the output. This checking is particularly important for a new, unproven pipeline. GenAI was trained to create believable output data, so ensuring that the accuracy of this generated data is within an acceptable threshold for the project is vital.

The Right Combination

Optimal outcomes require a combination of AI capabilities and human expertise. Consider the process of merging data from multiple sources, a common challenge due to overlapping names, birthdates, or other identifying information. GenAI can analyze additional data fields such as addresses, policy details, or claim histories to estimate the probability of correct matches between overlapping records. Actuaries can then calculate an overall risk score based on these match probabilities, improving accuracy in merged datasets.

To ensure actuarial soundness, actuaries must establish controls for potential errors introduced by probabilistic methods. They should implement validation procedures such as sensitivity analyses, confidence intervals, and cross-validation techniques to measure and manage these risks effectively. This structured approach enables actuaries to leverage GenAI's insights while maintaining rigorous standards of data quality and reliability.

Despite clear benefits, insurers must approach AI-driven data-cleaning responsibly. Predictive and generative techniques introduce ethical and regulatory considerations, particularly regarding transparency and auditability. Actuaries should implement robust validation processes, ensuring AI-generated outcomes are documented, explainable, and subject to routine audits. Clear communication of probabilistic data quality improvements versus exact deterministic fixes is essential.

Conclusion

GenAI and predictive AI present exciting opportunities to enhance actuarial data-cleaning practices. Actuaries who embrace these advanced AI techniques can achieve substantial improvements in efficiency, accuracy, and overall decision-making. By combining AI-generated insights with professional actuarial judgment, insurers can harness the full potential of their data assets, improving both business outcomes and strategic agility.

Jeff Heaton, Ph.D., Vice President, AI Innovation, Reinsurance Group of America, Incorporated (RGA)



**ACTUARIAL
INTELLIGENCE
BULLETIN**

AI in the Determination of the Actuarial Discount Rate

ANDREA MENTE

In actuarial evaluations of post-employment benefits, the discount rate is one of the most critical and technically sensitive elements of the entire measurement process. Its role in bringing future cash flows to present value produces significant impact on the actuarial liability. In practice, the choice of this rate is subject to scrutiny by auditors and plan sponsors—not only due to its financial impact, but also because it must be consistent with market reality and the prevailing regulatory principles.

In Brazil, in accordance with accounting standard CPC 33(R1), which is aligned with the international standard IAS 19, the discount rate must reflect, as of the balance sheet date, the market yields of high-quality credit instruments with maturities consistent with the benefit payment schedules. In countries such as the United States and European Union members, this guideline is implemented through yield curves based on highly rated corporate bonds (AA or above), issued by companies with ample liquidity in the secondary market.

Here in Brazil, however, there is no corporate bond market sufficiently deep and liquid to meet these requirements. The consolidated and accepted practice is the exclusive use of federal government bonds—primarily those indexed to the *Índice Nacional de Preços ao Consumidor Amplo*, or the Extended National Consumer Price Index (IPCA)—as the legitimate reference for the discount rate.

This regulatory and structural limitation presents a significant challenge: how to enhance the technical quality of the discount rate estimate in a context where the database is limited and the market lacks credit risk comparable to that of long-term actuarial obligations? This issue has been the subject of discussion among actuaries and plan sponsors. In this context, artificial intelligence (AI) emerges as a potentially transformative tool for future practice—if market conditions evolve to support its effective application.

Currently, the application of AI to the determination of the discount rate in Brazilian actuarial valuations is unfeasible, as there is no regulatory freedom to use curves derived from instruments other than federal government bonds. Still, the discussion remains relevant. Should the national financial market evolve—with a more robust corporate segment, regular issuances, consistent ratings, and adequate liquidity—conditions would emerge for AI to become an effective ally in estimating this parameter.

In developed markets, AI is already applied with increasing sophistication. In the United States, large actuarial consulting firms employ machine learning to calibrate discount curves based on AA-rated corporate bonds. These curves are derived from secondary market data and are adjusted using criteria such as outlier exclusion, liquidity correction, and spline interpolation. Supervised models are used to identify patterns between credit spreads and variables such as inflation, benchmark interest rates, and sector indicators.

In Europe, EIOPA adopts a framework based on the risk-free rate adjusted for volatility and extrapolated using the Ultimate Forward Rate (UFR). Local consultancies apply latent factor decomposition algorithms and unsupervised learning to model the hidden components of interest rate curves. The high granularity of available data and integration with regulatory models favor the use of AI not only for curve projection but also for incorporating complex economic expectations and scenario-dependent variations.

In the Brazilian market, the current regulatory framework imposes rigid constraints on the actuary's activities. Although it would be technically desirable to estimate discount rates based on corporate curves, such practice is still not used due to the absence of a market that permits statistically consistent inferences from observable data. Any attempt to use private bonds in the context of post-employment benefit valuations remains subject to audit rejection.

However, this reality may change. The maturing of the capital market—driven by regulatory reforms and the growing demand for long-term financing outside the banking system—tends to foster greater diversity and depth in corporate debt instruments. In this prospective scenario, the use of curves based on private bonds may become both technically defensible and normatively viable as a primary reference for the actuarial discount rate.

Once this stage is reached, artificial intelligence will find fertile ground for broad application. Supervised models such as gradient boosting and deep neural networks could be trained with historical data on corporate issuances, sector spreads, and macroeconomic assumptions to project dynamic discount curves with greater responsiveness to changing conditions. Time series techniques with recurrent neural networks, such as Long Short-Term Memory (LSTM), would allow for more stable estimates, less sensitive to short-term noise.

In addition, generative models—such as Generative Adversarial Networks (GANs)—could simulate interest rate curves under multiple macroeconomic scenarios, enabling probabilistic analyses of the discount rate. This approach would substantially enhance the quality of sensitivity analyses of actuarial obligations. AI could also detect statistical

distortions in databases, eliminate noise, and improve interpolation between maturities—central elements for building a reliable term structure.

Although regulatory restrictions currently prevent the direct adoption of these methods in Brazil, it is essential for the actuarial community to anticipate this discussion. Technical preparation must precede regulatory feasibility. Actuaries proficient in data science, multivariate statistics, and predictive modeling will be better positioned to lead the methodological transition that must accompany the structural advancement of the national market.

I conclude with a professional conviction refined through practice: the accuracy of the discount rate does not depend solely on the source of the data, but also on how that data is processed. When Brazil has a corporate credit market capable of properly reflecting the long-term risks of pension liabilities, artificial intelligence may serve as the bridge between empirical data and the construction of technically refined actuarial assumptions—in harmony with the principles of consistency, prudence, and neutrality that guide the actuarial profession.

Andrea Mente, FIBA Actuarial Partner ASSISTANTS Consultoria



AI Insights Meeting May 2025

On May 14 and 15, the SOA hosted the AI Insights virtual conference that featured 13 sessions over the two days with 33 speakers. This is the 2nd annual event and proved to be a great follow-up to the 2024 Inaugural conference, which is now available on PD Edge+. Here are summary descriptions of a sampling of those sessions.

Synthetic Data—Andrew Dilworth

In our session, my co-presenters (Tom Callahan and Joe Long) and I provided a high-level overview of synthetic data. We covered what it is, potential use cases, and ways to generate and evaluate it.

Synthetic data is artificially generated data that resembles real-world data. We highlighted advantages like increased volume, enhanced security, and reduced storage footprint, while also noting challenges such as ensuring realism, computational costs, and regulatory/professional compliance. We reviewed several methods to generate tabular synthetic data (rules-based, copulas, VAEs, and GANs) as well as sequential or time series data (LSTM, transformers). We discussed evaluation criteria focusing on fidelity, utility, and privacy. Finally, we concluded with a case study on generating synthetic health insurance claims data, demonstrating a practical application and lessons learned.

Federated Learning—Zhiyu Quan

In this session William Konoop and Frank Quan explore the critical balance between leveraging proprietary data for collaborative gains and protecting business-sensitive information and privacy. Taking advantage of unique and comprehensive proprietary data from multiple insurers and InsurTech firms, we present the first comprehensive examination of the economic benefits of the Federated Learning framework as a privacy-enhancing practical solution for data collaboration. We explore a complete variety of dimensions that span the boundaries of multiple business entities within the business insurance context. The proposed business solution could foster access to the collective intelligence of industry peers in scenarios where physical data aggregation is infeasible, thereby yielding a more comprehensive assessment of user-related features across entities and providing industry-level insights,

enhancing standards, improving efficiency, and driving innovation. We received many questions during the presentation and believe we triggered interest from the audience. Contact Frank Quan (zquan@illinois.edu) for a potential federated consortium, identifying industry trends in areas such as healthcare and mortality, or collaborating to combat fraud.

How to Use AI to Help you Learn—Effie Kalloo

Igor Nikitin and Dean Rootenberg outline Artificial Intelligence (AI) tools that can help actuaries work faster and more efficiently. AI is highlighted as the perfect sidekick.

A live demonstration of Notebook LLM by Google showcases its capability to turn any document into an engaging podcast. It can also create summaries, timelines, study guides and sample quizzes, which can be helpful for actuarial students. Similarly, another AI tool, Microsoft Visual Studio, integrated with GitHub can optimize, debug and explain program code functionality. Overall, this presentation illustrates the potential for AI to make the learning process more dynamic and efficient.

Lean on Me: Actuarial Professionalism and AI—Jon Wu

With an eye on the significant developments in artificial intelligence (AI) applications in the insurance industry in automation and underwriting, enhanced fraud detection, evolving risk modeling techniques, and the utilization of generative AI, Jason Reed, Mitch Stephenson and Matthew Wininger address emerging operational challenges, regulatory concerns, compliance risks, and model and data-related risks.

Health insurers' AI-driven decision-making have resulted in significant legal challenges and Zillow's \$500 million loss due to misvalued property assessments. These cases demonstrate the importance of establishing AI core principles, which encompasses minimizing bias in training data and outputs, establishing clear accountability through dedicated oversight, protecting privacy and confidential information, and maintaining transparency via comprehensive documentation and communication.

Meanwhile, Actuarial Standards of Practices (ASOPs) such as ASOPs 12, 23, 41, 54, and 56 that reinforce risk classification, data integrity, professional communication and reporting, and model reliability can help insurers to establish best practices. Actuaries are uniquely positioned to lead an ethical adoption of AI due to their professionalism and expertise in risk management by contributing to and keeping pace with AI advancements.

How to Get Comfortable with LLMs—Alex Zaidin

Large Language Models (LLMs) are rapidly transforming the insurance landscape, offering unprecedented opportunities for operational efficiency and innovation. As part of the broader AI ecosystem, LLMs enable life insurers to streamline complex processes across actuarial functions. However, successful implementation requires a strategic approach that balances technological potential with robust risk management.

Presenters Drew Caulk, Andy Rallis and Alex Zaidlin discussed key considerations around LLM risk management, including data privacy, maintaining model transparency, managing potential biases, and establishing rigorous governance frameworks. They touched on why actuaries are uniquely positioned to lead this technological revolution, leveraging strong quantitative backgrounds, domain expertise, and risk management mindset. They concluded that by adopting a measured, risk-aware approach, actuaries can harness LLMs to enhance decision-making, improve customer experiences, and gain competitive advantages while maintaining regulatory compliance.

Prompt Clinic—Dave Ingram

The focus was on practical techniques for improving large language model (LLM) prompting. Dave Ingram emphasized that prompting is not casual conversation but an act of interface design. Six key topics were presented:

specificity in prompts, defining the audience, understanding how the model generates answers, validating outputs to detect hallucinations, using chain prompts for complex tasks, and keeping a human (or actuary) in the loop. Ingram demonstrated prompt engineering with both single-shot and multi-step approaches, explored validation through questioning and counterarguments, and encouraged using outline-first strategies. He highlighted risks of LLM biases and overconfidence and stressed actuaries' responsibility for vetting AI-generated content. The session included live demos, active Q&A, and guidance on secure use of AI with sensitive data.

Practicing Actuarial Leadership with an AI Simulation—Dave Ingram

How can actuaries use AI-driven simulations to develop leadership and decision-making skills in complex, real-world situations? The simulation places participants in the role of Chief Actuary at an insurer facing potential bias in an AI-assisted hiring system. In a live demo Michael Niemerg interacted with AI personas—staff members, an HR leader, and a feedback expert—across four scenes, practicing communication, persuasion, and strategic thinking. The modular design of the simulation allows for easy adaptation: problem scenarios, counterpart personas, and feedback personas can be swapped to target different learning objectives and problems. Dave Ingram also shared the system prompts and structured role files. Simulations can provide realistic leadership practice, are easily customizable, require no coding, and can be deployed today using accessible AI tools.

Recordings of these sessions will all be available on PD Edge+ in the fall.



ACTUARIAL INTELLIGENCE BULLETIN

Actuarial AI Links

Artificial Intelligence Research from the SOA Research Institute

<https://www.soa.org/research/topics/artificial-intelligence-topic-landing/>

Artificial intelligence & machine learning from the SOA Emerging Topics Community

<https://www.soa.org/digital-publishing-platform/emerging-topics/>

Artificial Intelligence Research Listserv from the SOA Research Institute

<https://soa.wufoo.com/forms/join-ai-research-listserv>

American Academy of Actuaries AI Page

<https://actuary.org/topic/artificial-intelligence/> AI Tools for Actuaries
<https://aitools4actuaries.com/>

Actuarial AI Case Studies

<https://github.com/IAA-AITF/Actuarial-AI-Case-Studies>

AI-Insights Podcasts from SOA Research Institute

<https://getpluggedin.libsyn.com/>

AI for Actuaries from the IAA AI Task Force

<https://www.aiforactuaries.org/>



ACTUARIAL INTELLIGENCE BULLETIN

Are You a Typist or an Actuary?

STEPHEN J. MILDENHALL

I was dinged in my first performance review (in 1992) for not using the typing pool.

For younger readers: a typing pool was a group of secretaries whose job was to type up handwritten memos. I had taken touch typing lessons in high school, and I'd written my PhD thesis on a word processor, so I typed my own memos.

That, apparently, was not acceptable.

But the very technology that let me type my own memos - the personal computer - went on to transform the work of actuaries more broadly.

As PCs entered the workplace, they supercharged what actuaries could do. With spreadsheets, databases, and high-quality printed and on-screen graphics, it became possible to explore data interactively, visualize complex results, and iterate quickly. No more waiting overnight for a mainframe job to finish. The entire workflow changed - and actuaries became more powerful and more valuable as a result.

Now, with the arrival of large language models, we're standing at a similar inflection point.

As you assess the impact of LLMs on your profession, you need to ask: will your role disappear like the typist, or flourish like the actuary?

The key difference is domain expertise.

Typing pools were all process: no deep judgment, no specialized insight. So, when the technology came, it replaced them.

Actuaries brought domain expertise. Instead of being replaced, their skills were amplified. Over the past 30 years, demand for actuaries has grown dramatically. The number of qualified actuaries has increased more than five-fold since 1992, because their knowledge could be scaled through better tools.

LLMs are today's PCs. If your work draws on deep knowledge, analytical thinking, and professional judgment, LLMs will enhance your impact. But if your role is mostly about routine output, without real understanding behind it, you may be in typing pool territory.

History doesn't repeat, but it often upgrades: which path is your profession on?

Stephen J Mildenhall, PhD, ASA, FCAS, MAAA retired blogger at <https://blog.mynl.com/>

A dark blue banner with a light blue circuit-like pattern of lines and dots. The text "ACTUARIAL INTELLIGENCE BULLETIN" is written in white, bold, uppercase letters on the left side.

ACTUARIAL
INTELLIGENCE
BULLETIN

Editorial Committee

Jon Forster, ASA
Dave Ingram, FSA
Ronald Poon Affat, FSA
Frank Quan, PhD
Darryl Wagner, FSA

Associate Editors

Jing Kai Ong, ASA
Jeremy Levitt, FSA

Thank you for reading the 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 rpoonaffat@soa.org.

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