

Reserving Judgement

An answer to a question about AI and the actuarial career

AI-use disclosure

This manuscript was prepared with assistance from Claude (Anthropic) and OpenAI Codex. All AI-assisted output was critically reviewed, verified against source material where relevant, and substantially revised by the author. The selection and interpretation of evidence, analysis, arguments and conclusions are the author's own. The author accepts full responsibility for the final manuscript.

Originality and publication

I affirm that this essay is my original work and has not previously been published. I grant the Actuarial Society of South Africa a non-exclusive, perpetual licence to publish and distribute it with appropriate attribution.

Opinions

The opinions expressed in this essay are those of the author and do not necessarily reflect the views of the Actuarial Society of South Africa. They also do not necessarily reflect the views of EY.

1. The student's question

At a graduate recruitment drive at a South African university, an honors student asked me a question I have not stopped thinking about since: With AI changing actuarial work, what should I be doing now to prepare??

The common answer is to start learning AI. Use the tools, understand what they can and cannot do, the risks they pose, and become comfortable working with them. That answer is right, but it is too limited. She was asking not only how to prepare for the first day of her first job next year, but how to prepare for the forty years after it, during which the tools and the work will change repeatedly.

Much has been written about what AI will do to actuarial work. The deeper question, and the subject of this essay, is what working with AI may do to actuaries: what they learn, how they think, what they understand and how they take responsibility. This essay is my attempt at that larger answer.

I am reminded of a quote from John Culkin in 1967: "We shape our tools, and [then] our tools shape us." Culkin was writing about media technologies, television, film, radio and the new electronic media of his time, and his point was that they did more than transmit information: they trained people to notice differently, think differently and organise classrooms, politics and public life differently. AI will do the same to actuaries, not only by changing the speed of the work, but by changing what juniors practise, what seniors trust and perhaps even how the profession defines itself.

Consider what the spreadsheet did. Before it, a valuation was carried out against printed tables in bound books, by hand, and the binding constraint on the work was arithmetic. The spreadsheet moved the arithmetic into the machine and changed how long a valuation took, how many bases could be tested, how fast an error could travel and how many people had to be in the room. It did not change who was building the model; actuaries still designed the model and the machine still only calculated the result.

Even so, the spreadsheet reshaped the actuary from the outside: where we worked, how fast, against which tables, with what records. AI stretches further because it is crossing the line from instrument to participant. A spreadsheet waits for an actuary to specify the calculation; an agentic system can be given an objective, decide which steps to take, operate the files and workflows in which the work lives, test what it finds and return when it needs a judgement (or even make the judgement call itself in certain cases). The consequential change is not that the tool can answer us, but that it can begin to think and act alongside us or on our behalf.

So what this tool reshapes is not only the work. It is the professional: how you learn, how you think, how you take responsibility and how you see yourself.

2. What nobody can tell you

Nobody knows exactly what actuarial work will look like in ten years. Forecasts have failed in both directions. An IFoA survey illustrates the range of expectations: 44% anticipated significant changes in actuarial roles and responsibilities, 41% expected some change while core roles remained, and 12% expected complete transformation.

What we can observe is more specific and more useful. AI is currently able to automate tasks, not yet entire jobs. Some routine production tasks shrink effort, some work becomes reviewing what AI produces, and new work using AI emerges around oversight, evaluation and workflow design. Actuarial roles are recomposed as the balance between production, interpretation, judgement and accountability changes.

This does not necessarily make the work easier for actuaries. AI may reduce the effort required for each task while increasing the pace and difficulty of the job. Faster production means work circulates across the organization faster, and productivity may be absorbed as greater throughput rather than reduced work time. As more of the routine work is automated, the remaining human contribution will likely become more and more about judgement, validation and responsibility. The human may actually become the bottleneck in a faster pipeline.

The consequences may first impact new graduates at the entrance to the profession. If routine work shrinks, juniors may face a narrower entry point and be required to exercise judgement and take accountability much earlier.

3. The most exposed generation

Actuarial work can be understood as four overlapping layers: production, interpretation, judgement and accountability. AI pressure is strongest near the bottom of this stack, but it is moving unevenly upwards.

Production is where the work product is made, moved, formatted and checked: data preparation, model runs, exhibits, documentation and supporting workings. It is the most digital and repetitive layer, and therefore the most immediately exposed. As speed becomes cheaper, human value shifts from performing the work to controlling and validating it.

Interpretation connects the output to the business: reconciling results,

analysing movements, interpreting stresses and drafting commentary. This layer is more exposed than it appears because a fluent explanation is now easy to produce. The important human question is no longer whether the explanation reads well, but whether it is true, complete and material.

Judgement begins when somebody must decide what is reasonable, credible or worth escalating. Certain judgement in certain circumstances is likely at risk. It is not the protected category actuaries sometimes imagine. Years of reports contain repeated decisions, objections and thresholds from which a system can infer how similar situations were judged before. Live context, what matters now, for this business, under these conditions, remains harder.

Accountability is different. AI may assemble the evidence, test the reasoning and recommend a conclusion, but it cannot occupy a professional office or bear the consequences of being wrong. A statutory actuary can be challenged, sanctioned and replaced. Accountability is durable not because it is necessarily the most difficult layer, but because it is institutionally assigned. Capability does not transfer duty.

The boundaries between these layers can be blurred, and most actuarial work contains some of all four. The categorization is useful because of the impact from AI.

Juniors spend more of their time in production and interpretation; seniors carry more judgement and accountability. AI will affect both, but it lands most directly on juniors. Evidence outside the actuarial profession already suggests that labour-market adjustment may appear first at this entry point. Stanford's Digital Economy Lab found employment among workers aged 22 to 25 in highly AI-exposed occupations about 19% below where it would have been had it tracked their less-exposed peers, with the difference concentrated in hiring. The finding is descriptive rather than causal and does not predict actuarial employment, but it shows where pressure may emerge first.

There is also a generational asymmetry. Senior actuaries built their mental models by doing the underlying work: constructing models, chasing movements, writing weak reports and being challenged. They now place AI on top of expertise that already exists. Juniors may use the same tool before that expertise has formed; meaning they have to simultaneously develop their skills in using AI and develop the judgement and intuition that their more senior colleagues did.

Junior work has always produced two things in the same hours: the work product and professional experience. AI can improve the first while weakening the second.

The risk now is not losing your job to AI, but losing the work that would develop

you into a more capable senior actuary.

4. The four separations

AI impacts your development by pulling apart four things your seniors could take for granted as fused. Each separation follows a similar pattern: something that was one thing becomes two, the work can still be done, and the person is no longer made by doing it. In each case there is a discipline that puts the two back together deliberately, and in each case the discipline costs something.

4.1 *Doing from learning*

Apprenticeship use to be embedded in doing the work. Rebuilding a model from source data taught you structure. Writing a weak assumption note taught you the difference between convention and judgement. Peer reviewing a colleague taught you which errors recur and their implications.

Routine work did more than simply produce an output. Its repetition gave juniors a baseline; many iterations and reviews gradually turned into judgement and intuition. AI can now produce the output while robbing the analyst actuary of those formative experiences.

This is difficult to avoid because each delegation to AI seems rational and is incentivized. The work is done sooner, with less effort. The manager receives the work sooner. The stakeholders get a shorter turnaround. The client ultimately pays less and the firm improves its margins. Learning is eroded in the pursuit of efficiency.

Delivery work and formation work may now have to become separate. In delivery work actuaries should use the machine aggressively wherever it is competent, permitted and governable; a good opportunity, in fact, to practice sought-after AI skills. In formation work, actuaries must deliberately preserve the learning that the delivery work no longer provides: build a simplified model from a blank workbook; estimate the answer before you ask the model; explain the result with the report closed; reconstruct generated code and defend its assumptions.

AI can be a part of that formation work but as a tutor rather than a copilot. It should ask questions, demand explanations and make you do the work of arriving at the answer.

The uncomfortable truth is that, while you are learning, you may sometimes be slower than peers who optimize for delivery. That slowness can be a worthwhile investment when it preserves the practice from which deeper judgement develops.

4.2 Writing from thinking

For difficult work, writing is never the transcription of finished thought; the draft is where the thinking happens. George Orwell famously said, "But if thought corrupts language, language can also corrupt thought." And anyone who has written a complex actuarial report would certainly agree.

You begin believing the deterioration in the company's claims is mostly exposure-driven. Then you try to write the paragraph, and the historic loss ratios don't quite fit the sentence you've written, so you go back to the data and find the business mix has shifted. The qualification you add then reveals that another assumption is inconsistent. The rewritten pages were not waste; they were the mechanism that produced the surviving insight.

AI can provide the polished paragraph without the analyst having to think through the issue. Much of the thought process becomes optional. The risk is not bad writing, it is a well written report unsupported by understanding.

I should clarify, not all writing is about thinking. Transcription, formatting and rephrasing can be delegated without loss. What we should protect is epistemic writing, undertaken to discover what you think. Outsource the difficult first draft and you outsource the reasoning it would have forced.

Write the difficult first drafts yourself. Use the machine socratically (using it to question you): make it question your explanations, challenge your claims or draft in parallel so that you can understand the difference.

The discipline is not to write every word yourself. It is to ensure the conclusion has been thought before the machine makes it sound finished.

4.3 Output from understanding

You generally understand a difficult output largely because producing it requires understanding. AI makes it possible to possess the output without possessing the reasoning.

The danger is that you may not notice the difference. A fluent explanation feels obvious once you have read it, and recognition is easily mistaken for understanding. "That makes sense" is not the same as "I know why this is true."

Understanding is generative. You can reconstruct the reasoning, connect the conclusion to the evidence, explain how it would change under different assumptions and identify where it might fail.

So test your understanding before accepting AI-assisted work. Form a provisional view before asking for the model's. Then put its answer away and explain the conclusion in your own words. Trace each material claim back to the source. Reproduce the critical calculation or logical step. State the

strongest alternative explanation and the evidence that would change your mind.

Use the machine as an examiner as well as an answerer: ask it to challenge your assumptions, vary the facts and question you without supplying the answer.

A useful test is whether you can defend the work once the output is no longer in front of you. If you cannot yet reconstruct the reasoning, connect it to the evidence and explain where it might fail, the polish of the output should not be mistaken for understanding.

4.4 Capability from accountability

The first three separations culminate here. AI may help you perform, explain and understand work beyond your experience. The remaining question is what responsibility you can take for it.

When a senior delegates to a junior, the work moves through a chain of human accountability. The senior remains responsible for supervision and sign-off. The junior actuary becomes professionally responsible for the quality of their contribution, for working within their competence and for making uncertainties and limitations visible as the work moves upwards.

When you delegate to AI, the task moves outside that chain. The system can perform the analysis, but it cannot accept a professional duty. It may identify a limitation, but you cannot rely on it to recognise that the limitation is material, communicate it to the right person or ensure that somebody acts.

AI therefore transfers execution without adding another accountable participant. The professional buck stops with the human using it.

For a junior actuary, this may bring accountability as a skill forward in your career. You may have to take responsibility for work you did not personally produce very early on in your career. This does not mean assuming authority you do not have. It means knowing what you are claiming, what you have relied on, what remains uncertain and who needs to know.

Depending on the situation, accountability may require further testing, disclosure of material AI use, correction of an error or escalation to somebody with greater competence or authority. The principle is that responsibility must remain visible and assumed by a person.

AI can extend your capability beyond your experience. Becoming a professional means learning what you may own, what you must disclose and what you have a duty to pass upwards.

5. The new professional skill: allocating intelligence

The four separations will not manage themselves. They create a new professional skill: allocating intelligence wisely which means deciding what you should think through, what the machine should do, where the two should collaborate and who must own the result.

Wise allocation begins with purpose. Are you using AI for speed, breadth, challenge or learning? Then ask two questions: what happens to the work if AI does this, and what happens to me? The first concerns quality, verifiability, materiality and accountability. The second concerns the learning, thinking and understanding that delegation may remove.

A routine, reversible and readily checked task is a reasonable candidate for outsourcing. Exploration may benefit from collaboration, provided you retain an independent view and keep the reasoning traceable. Material judgement that is difficult to verify should remain human-led, with the machine used as a counterparty to attack the conclusion, generate alternatives and expose missing evidence.

The answer should also depend on the actuary. The same delegation can be leverage for an expert and lost practice for a trainee. Sometimes the right choice is speed; sometimes it is deliberate effort; sometimes it is asking for challenge rather than an answer. Maximum automation is not the objective. The objective is better work without surrendering the human capability required to judge and own it.

Most importantly, recognise you are not adopting a stable tool; you are adapting to a moving capability frontier. The models will change, the workflows around them will change and your own competence will change. A task you should perform yourself while learning may later be sensible to delegate. Work the machine could not do last year may become routine next year. Every allocation decision therefore has an expiry date.

Adoption is an event but adaptation is a professional habit. It means repeatedly testing what the system can do, where it now fails, whether the controls still work and what its use is doing to your own capability.

Machine fluency is therefore not mastery of today's prompts or interface. It is the ability to keep recalibrating the boundary between the machine's work and yours. Domain depth tells you what matters, what is plausible and where the risks lie. Machine fluency helps you act on that judgement as the frontier moves.

The new professional skill is not knowing what AI can do nor is it being skilled at how to use it, though both of these are important. The new professional skill is

being skilled and using it wisely and responsibly: deciding, again and again, what it should do, what you must still do and who must own the result.

8. The final judgement

Choosing not to use AI is still possible today, it is unlikely to remain so. Today an actuary may need to justify using AI. As these systems become better than even an experienced actuary at a growing range of actuarial tasks, there will be cases in which the harder question is why they were not used. If a secure and validated AI-assisted review consistently finds errors that manual review misses, a human-only process is not more professional merely because it is human. Using AI may become a professional and ethical responsibility in certain cases.

But it is likely that accountability will not follow this path. The machine may become the best-performing part of the reasoning chain in an increasing number of areas of actuarial practice; professional responsibility is likely to remain human.

So use AI regularly, and keep relearning it as the frontier moves – and it is moving quickly. Delegate routine, reversible and readily checked work; preserve the formative effort that is still helping you to develop skills. This will determine quality of professional experience you will draw upon years later. Do not outsource your thinking to AI; there is no substitute for epistemic writing. Form a view before asking AI in situations when judgement matters. When AI gives you an answer, put it away and prove that you can reconstruct the reasoning. Make material reliance and uncertainty visible, and escalate what exceeds your competence or authority. Use the machine not only as an answerer but as a challenger.

At that recruitment drive, the student asked what she should do with AI. Culkin's observation reveals the deeper question: what will her repeated use of AI do to her and her career trajectory? Every prompt, workflow and control helps shape the tool's place in the profession; every choice about what to delegate shapes the professional in turn.

We shape our tools, and then our tools shape us. Use the tool. Learn from it. Challenge it. But be intentional about those choices, and alert to their impact on you and your work.

That is one judgement you should never delegate.

References

- Actuarial Society of South Africa (ASSA) (n.d.) *Professional resources and conduct*. Available at: [ASSA professional resources and conduct](#) (Accessed: 1 September 2026).
- Actuarial Society of South Africa (ASSA) (2022) *APN 106/403: Head of Actuarial Function for South African insurers*. Version 7, effective 1 December 2022. Available at: [ASSA APN 106/403](#) (Accessed: 1 September 2026).
- Brynjolfsson, E., Chandar, B. and Chen, R. (2026) *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Revised 12 August. Stanford Digital Economy Lab. Available at: [Stanford Digital Economy Lab working paper](#) (Accessed: 1 September 2026).
- Cisco (n.d.) *What is agentic AI?* Available at: [Cisco: What is agentic AI?](#) (Accessed: 1 September 2026).
- Culkin, J.M. (1967) 'A schoolman's guide to Marshall McLuhan', *Saturday Review*, 18 March, pp. 51–53, 70–72. Available at: [A Schoolman's Guide to Marshall McLuhan](#) (Accessed: 1 September 2026).
- Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rostaniec, K. and Troszyński, M. (2025) *Generative AI and jobs: A refined global index of occupational exposure*. ILO Working Paper 140. Geneva: International Labour Office. Available at: [ILO research repository](#). <https://doi.org/10.54394/HETP0387> (Accessed: 1 September 2026).
- Institute and Faculty of Actuaries (2025) *Ethical and professional guidance on data science and artificial intelligence: A guide for members*. Version 2.1, July. Available at: [IFoA ethical and professional guidance](#) (Accessed: 1 September 2026).
- Institute and Faculty of Actuaries (n.d.) *Standards and guidance*. Available at: [IFoA standards and guidance](#) (Accessed: 1 September 2026).
- Institute and Faculty of Actuaries and Financial Reporting Council (2025) *Joint IFoA and FRC regulatory update: Guidance relevant to AI, data science, and machine learning*, 26 March. Available at: [Joint IFoA and FRC regulatory update](#) (Accessed: 1 September 2026).
- Microsoft (2026) *Agents, human agency, and the opportunity for every organization*. 2026 Work Trend Index Annual Report, 5 May. Available at: [Microsoft 2026 Work Trend Index](#) (Accessed: 1 September 2026).
- Orwell, G. (1946) 'Politics and the English language', *Horizon*, 13(76), pp. 252–265. Available at: [The Orwell Foundation](#) (Accessed: 1 September 2026).
- Society of Actuaries (n.d.) *Actuarial tables, calculators and modelling tools*. Available at: [Society of Actuaries tables and tools](#) (Accessed: 1 September 2026).
- Zhu, B., Ramsay, D.C. and Dagli, D. (2025) *Generative AI in the actuarial profession: Survey insights from the IFoA GenAI Working Party—transformation, adoption, and the road ahead*, 22 May. Available at: [IFoA GenAI Working Party survey](#) (Accessed: 1 September 2026).