





AI Won't Make You Dumber. It'll Do Something Worse.



Frank Barilone Jr., MBA, CSM

July 14, 2026

I've built executive decks for twenty years. For most of that time, the deck was a side effect. To get to a slide you had to reconcile a number, discover the number was wrong, chase down why, and decide what it meant. The slide was the residue of that process. What you actually walked away with was a working model of the problem in your head, and you could defend any line on any page because you had personally earned every one of them.

Last month I built a fourteen-slide strategy deck with AI doing most of the assembly. It was better than what I used to produce. Cleaner structure, tighter narrative, a fraction of the

time. I also noticed something. The parts I reconciled myself, I still knew cold. The parts I didn't, I knew the way you know a movie you watched once.

That gap is the thing nobody is talking about yet.

The proxy came unwelded

Knowledge work has always measured the artifact, because the artifact is visible and thinking isn't. The memo. The deck. The model. The analysis. We rewarded output because output was a reliable proxy for understanding, and it was reliable for exactly one reason: you could not produce the artifact without doing the thinking. The proxy was load-bearing because it was welded to the thing it stood for.

AI unwelds them.

You can now produce the artifact without the understanding. And because the organization only ever measured the artifact, it won't notice. It will promote you. You'll ship four decks instead of one and look like the most productive person in the room, right up until someone asks a question that isn't on a slide.

This is Goodhart's law arriving in knowledge work with no defenses in place. The measure became the target the moment the measure stopped requiring the thing it measured.

Effort was never the variable

The lazy version of this argument is that hard work builds character. That's wrong, and it's worth saying so plainly, because it's the version that will get this whole idea dismissed.

Three hours aligning text boxes builds nothing. Formatting a table of contents builds nothing. Effort is a terrible proxy for understanding, and if your defense of manual work is that it was hard, you deserve to lose the argument.

What builds understanding is ambiguity resolution. Judgment forms in the specific moments where the work forced you to make a call under uncertainty and then own it. The number doesn't tie and you have to decide which source is right. Two stakeholders want incompatible things and you have to pick. The data supports three stories and you have to choose one and be accountable for the choice.

AI strips effort out of formatting and out of judgment with equal indifference. It doesn't know the difference. You have to know the difference, and you have to keep the second one on purpose.

Most people aren't. The temptation is to hand over the entire task, because handing over the entire task is what feels like leverage.

We are spending down an inventory

Here's why senior people are slow to see this coming.

I can let a model build the deck and still catch the bad number, because I accrued twenty years of pattern recognition in an era where there was no alternative to accruing it. I'm not losing expertise. I'm just no longer adding to it.

That's survivable for me. It is not survivable for the profession.

Judgment gets built by junior people doing unglamorous work badly, getting corrected, and slowly compiling a model of how things actually go wrong. The grunt work was not a tax on becoming an expert. The grunt work was the training data. It's also the first thing AI ate.

So the damage doesn't show up next quarter. It shows up in one professional generation, when the senior who can smell a bad number was never manufactured, because nobody ever made her build the thing by hand.

Look at the timeline on social media. It arrived in 2004. Broad consensus that something had broken took until roughly 2020, and it required a cohort that had grown up inside it. Scale went up immediately. Intimacy went down slowly. Sixteen years and a full generation before the bill came due. The lag is the entire lesson, not the scale. A harm that takes a generation to surface is a harm no quarterly metric will ever catch.

The case that none of this matters

Now the part that should make you uncomfortable, because the optimist argument is stronger than most skeptics will admit.

If AI eventually supplies the judgment, then preserving human judgment is nostalgia, not risk management.

We don't require pilots to navigate by sextant. We don't require accountants to do long division. We let those capabilities go, and we were right to, because the machine got better than the human and we covered the failure modes with redundancy instead of with human backup. Nobody mourns them. Nobody should.

Run that forward. The junior analyst never develops the instinct to catch the bad number, and it doesn't matter, because the model catches it, and catches it more consistently than the analyst ever did. The erosion of human expertise is only a problem if human expertise stays load-bearing. Remove that assumption and my entire complaint collapses into sentiment.

I think that's a serious argument. I also think almost nobody has followed it to where it actually goes.

Follow it all the way

If the machine does the work and the machine supplies the judgment, what is the human for?

The honest answer is that the human is for liability.

Machines can't be sued. They can't be fined, sanctioned, debarred, or have a license pulled. They can't be deposed. They can't sign an attestation that means anything in a courtroom. Accountability requires a legal person, and for now that means a human being, or a company staffed by human beings.

Follow the automation curve to its end and the roles that survive are the ones where a human is required to *exist*, not required to think. The signer. The licensed one. The name on the attestation. The body on the regulatory org chart. Employment persists as a liability substrate rather than a cognitive contribution.

That is not a comfortable thing to write and I'm not going to soften it. Every knowledge worker should look hard at their own role and ask which parts of it survive that filter. For a lot of us, the answer is going to be narrower than we'd like.

The trap inside the liability moat

But hiding behind liability is not the safe harbor it looks like, and this is the part I want people in my industry to actually sit with.

A signature is supposed to mean somebody checked. That is its entire function. If the person signing has no capacity left to challenge what they're signing, the signature is a formality wearing the costume of a control. The paperwork says the control operated. It didn't.

This isn't hypothetical in mortgage servicing. Model risk management, third-party oversight, effective challenge: all of it rests on one premise, that somewhere in the process a human being understands enough to push back. If AI quietly erodes the judgment of the people whose job is to challenge AI's output, the control is hollow while every artifact documenting it insists the control is fine.

That is a governance failure that is invisible on paper by construction. And it's the one I'd expect to surface in an exam three years from now, on a decision made today, evaluated against a standard that didn't exist yet, by a regulator asking a question that was never on a slide, about a model that has since been retrained twice.

Aviation ran straight into this and paid for it in lives. Automation dependency and manual skill decay are documented, investigated, and understood. The industry's answer was not to remove the pilot, and it was not to switch off the autopilot. It was to push manual flying and recurrent training back into the seat, deliberately, so the human retains the capacity to take the controls on the day the system hands them back.

That's the model. Not resistance to AI. Deliberate, structural, expensive retention of the judgment that the work no longer builds for us as a byproduct.

The window is open now and it will close

Everything above says the loss is invisible to the organization. Sit with what that actually means for you.

If the org can't see it, the org will never require you to prevent it. No manager assigns this. No OKR captures it. No performance review has a line for it. The system will reward the person next to you who skipped it and shipped faster, and it will keep rewarding them for years.

So it's yours. Nobody is coming.

You also have a window. Somewhere ahead of us the machines absorb the judgment too, and this whole discipline becomes a hobby, like navigating by the stars. Maybe that's five years out. Maybe twenty. Nobody knows, and anyone who tells you they do is selling something. But between here and there is where your entire remaining career happens, and in that stretch judgment is still the whole game.

The room is your manual flying.

Nobody in your audience cares that a model built the slide. They're looking at you. Your name is on the deck. You are the one who has to take the question that isn't on the page, and the question that isn't on the page has always been the only question that mattered.

Stop treating the presentation as a delivery event. Treat it as a proficiency check. It's already on your calendar, it can't be outsourced, and it is the last honest test of whether you know what you claim to know, because the machine doesn't have to stand there and you do.

Hard rule: don't present what you can't defend cold. Not "can look up." Not "it's in the appendix." Defend, live, under pressure, when the number gets challenged by someone who's been doing this longer than you have. If you can't do that with page nine, you don't own page nine, and it does not matter who typed it.

Make the model train you as hard as you're training it.

Most people prompt for output. Prompt for transfer.

Make it write speaker notes, and don't treat them as presentation polish. Treat them as the model showing its reasoning back to you. That's the derivation. Read it like a student, not like an editor.

Make it state its assumptions out loud and flag where it's least confident. Then go verify exactly those things by hand. That's your reconciliation queue, and it's the highest-yield hour you'll spend all week.

Make it play the room. Ask for the ten hardest questions a skeptical executive would put to you, then answer them yourself, out loud, without looking. Every one you fumble is a gap you just found for free. Finding it in your office is a lot cheaper than finding it in the meeting.

Make it show the work, then spot-check one step at random. You're the auditor now. Behave like one.

Dial in everything that isn't judgment.

Lock your brand standards down to the hex codes. Fix the type stack, the slide architecture, the message hierarchy, the way you name things, the way you structure an argument. Not because it looks polished, though it will. Because every degree of freedom you leave open is a degree of freedom you have to review, and your review capacity is finite and it is now the scarcest thing you own.

If the model is free to vary the colors, you have to check the colors. Take that away from it. Constrain the machine to the surfaces that don't matter so the only variable left in front of you is the one that does. That's the whole trick. Push the machine's variance into formatting and pull your attention into substance.

Do it now, while your footprint is small. This compounds. Every unstandardized artifact you ship becomes a precedent somebody copies, and eighteen months from now you're not governing a process, you're excavating one.

If you have people under you, you owe them the same thing.

The work no longer builds judgment as a byproduct, which means somebody has to build it on purpose or watch it not get built. Make them do the hard part by hand. Make them defend it in front of you before they defend it in front of anyone else. If your juniors never reconcile a number that doesn't tie, you're not developing anyone. You're staffing.

Both futures reward the same discipline

This is what makes the cost worth paying either way.

If the machines never fully absorb judgment, the practitioner who kept theirs is the most valuable person in the building, because everyone around them traded it for speed and got promoted for it.

And if the machines do absorb it, and we all end up as signers holding liability we can no longer personally verify, then the only person qualified to hold that role is the one who still understands the system well enough for the signature to mean something.

The last person with judgment doesn't become obsolete. They become the control.

The trade may well be worth it. In most domains I think it will be. But make it with your eyes open, and be honest about what you're trading, which is your ability to know when the machine is wrong.

That capability used to come free with the work. It's about to become something you buy on purpose. Buy it. You won't find out you needed it until the day you needed it.

Frank Barilone Jr., MBA, is a product management leader with over 20 years of experience in consumer banking and lending, with a focus on strategic governance and AI adoption in regulated financial services. He is an active contributor to industry standards efforts in banking and finance.

frankbarilone.com