

ESSAY 11

BETTER JUDGMENT IN THE AGE OF AI

A Control Is Not Evidence That Control Works

"A signature is not a decision."

A Control Is Not Evidence That Control Works

By Alen Mayer

The advice about AI has settled into three words: keep a human in the loop. It isn't wrong. It's incomplete, and the missing part matters more than the part everyone repeats.

Much of the conversation about AI oversight eventually arrives at the same reassurance. A human reviews the recommendation before it becomes a decision. Say it enough times and it starts to sound like an answer. It isn't. It's a design choice, and design choices still have to prove they work.

■ A REVIEW PROCESS WORTH ENVYING

Consider a company that has built exactly this kind of control into its AI-assisted procurement process. Every recommendation the system produces requires a human sign-off before it moves forward. The quarterly numbers are as clean as numbers get: 100% of recommendations reviewed. Every review recorded, timestamped, signed. Not one missed step in three months.

On paper, there is nothing to discuss. The control exists. The control is followed. The control is documented. What more would there be to establish?

Then someone asks the reviewers a simple question: what do you actually look at before you sign?

■ WHAT COMPLETION PROVES, AND WHAT IT DOESN'T

When the reviewers explain their routine, a different picture emerges. Most read the AI's recommendation, find it reasonable, and approve it. Not because they've checked the assumptions underneath it. Because it arrives sounding like it's already been checked. The review happens. The scrutiny mostly doesn't.

That gap has a name: the Assurance Gap. The difference between demonstrating that a control was completed and establishing that it achieved what it was built to do. A human review can be performed, recorded and fully compliant with procedure, and still tell you far less than the paperwork suggests about whether anyone exercised judgment inside it.

A completed review is not proof that meaningful judgment occurred.

■ WHAT THE 100% FIGURE GETS USED FOR

The completion rate rarely stays a private number. It shows up in the board deck as evidence the AI system is safely governed. It gets cited when someone proposes expanding what the system is allowed to decide on its own, or loosening the review requirement because the process has "proven reliable."

That's where the Assurance Gap stops being a paperwork problem. The organization isn't just misreading its own controls. It's making its next decision — about authority, about scale, about how much oversight to remove — on confidence those controls never actually earned.

■ WHY THIS ISN'T AN ARGUMENT AGAINST CONTROLS

This isn't an argument against governance, against process, against the human-in-the-loop principle behind most AI oversight today. Reviews, audits, escalation steps and sign-offs still matter. Removing them doesn't close the Assurance Gap. It only removes the evidence you'd need to notice it.

The problem isn't that the control exists. It's that its existence gets treated as the finish line. A compliance report that says every review happened answers a narrower question than the one leadership actually needs answered: whether those reviews changed, challenged or meaningfully tested anything at all.

■ THE QUESTION THAT CLOSES THE GAP

There's no single fix here, and it would be too easy to say the answer is stricter controls or more reviewers. Sometimes it is. Sometimes it's better training, different evidence, more time, or a redesigned approval step. What determines which intervention is worth making is a question asked earlier: what evidence would show that this review is doing its job, beyond the fact that it happened?

For the procurement team in this example, answering it means going back to a sample of already-approved decisions. The team first identifies the material issues each sampled decision required a reviewer to examine, then compares those requirements with the evidence available from the completed review. Approval isn't the thing to look for — plenty of well-examined recommendations deserve approval. The question is whether the material issues in that decision were actually examined, not whether the reviewer changed the outcome.

A record of review is useful evidence. It just isn't sufficient evidence. It can tell you a step happened. It can't, by itself, tell you whether the judgment that step was designed to produce actually took place.

A signature is not a decision.

Companion Judgment Challenge: "The Control That Passed."

TAKE THIS FURTHER

Build better judgment into how your team works with **AI**.

I help leaders and organizations strengthen the human capabilities behind better judgment as AI becomes part of everyday work and decision-making.

JUDGMENT LABS

A 90-minute session working through realistic AI-influenced business decisions.

WORKSHOPS

Practical training on how teams question, interpret, weigh and decide with AI.

KEYNOTES

Reframing what human judgment requires as AI becomes more capable.

Alen Mayer

alenmayer.com