ESSAY 08

BETTER JUDGMENT IN THE AGE OF AI

# Who Thinks First Matters

*"AI doesn't only influence the answer. Sometimes it influences the thinking that happens before you reach one."*

**Alen Mayer**

alenmayer.com

# Who Thinks First Matters

By Alen Mayer

## One useful framework draws a distinction between two ways of working with AI.

In what it calls centaur mode, a person develops the thinking first and AI helps execute or elaborate it. In cyborg mode, AI takes the first pass and the person reviews, refines, and adds judgment on top. The argument is sensible: which mode fits depends on the task, the stakes, the person's own expertise, and what the AI is actually good at.

That's a real design choice, and it's worth making deliberately. But there's a question sitting underneath it that the distinction doesn't quite reach.

Does seeing AI's answer before forming your own change the judgment you're subsequently capable of making?

### THE SAME REVIEW, TWO DIFFERENT STARTING POINTS

A sales manager needs to understand why a major deal fell through.

In one version, she looks at the evidence herself first — the call notes, the email cadence, where the deal stalled — and forms a working view before asking anything of AI. Then she asks it to challenge that view: find evidence that contradicts it, surface alternatives she might have missed, argue the other side. The AI's job is to test a judgment that already exists.

In the other version, she opens the same evidence and immediately asks AI what happened. It replies with something coherent and plausible: price sensitivity, a weak executive sponsor, a competitor with a stronger technical answer. She then reviews the underlying evidence herself, checking whether it supports the explanation.

In both versions, she's the one who signs off. Both would show up identically on paper as "manager reviewed the analysis." But she isn't doing the same thing in each one. In the second, she isn't approaching the evidence to see what story it tells. She's approaching it already holding a story, checking whether the facts are roughly consistent with it — which is a much easier bar to clear than building the explanation from nothing, and a much easier way to miss whatever the first plausible answer didn't account for.

### WHY ORDER CAN CHANGE WHAT YOU SEE

This isn't a claim that AI's first-pass answer is usually wrong. Often it's a perfectly reasonable read, and reviewing it is a legitimate way to work. The issue is narrower:

once a plausible explanation exists, building a genuinely independent second one can become harder. The first answer doesn't necessarily sit there as one option among several. It can become the reference point against which everything that follows is interpreted.

> The first answer can become the reference point against which everything that follows is interpreted.

A postmortem team investigating a production outage runs into the same thing. Ask an AI system what caused an outage, and it will typically deliver something coherent within minutes — a specific service, a specific change, a specific mechanism. Once the team has that answer in hand, "reviewing" it can quietly become a search for whether the logs are consistent with the story, rather than a search for what else could produce the same symptoms — the activity has changed without anybody consciously deciding to change it. A team that instead built its own timeline and its own top three hypotheses before consulting AI at all is positioned to notice if the AI's answer, however fluent, is explaining only part of what happened. A team that saw the AI's answer first has a harder time noticing that, because the alternative now has to displace an explanation that already sounds coherent.

## THIS ISN'T AN ARGUMENT AGAINST LETTING AI GO FIRST

The instinct here could easily curdle into "always think it through yourself before you look at what AI says," and that would be both impractical and frequently wrong. For routine, well-understood problems, letting AI take the first pass is often exactly the efficient choice — there's no meaningful independent judgment being protected by insisting a person draft the standard email or the first version of a summary from scratch. The same framework's own examples make this point: some tasks genuinely suit AI-first execution, others genuinely suit human-first concept development, and the difference tracks the stakes and the ambiguity of the problem, not some fixed rule about which order is virtuous.

The deal-loss review and the outage postmortem are worth protecting from AI-first framing not because they're important in general, but because they're exactly the kind of problem where the first explanation anyone hears has an outsized chance of becoming the only one anybody seriously considers — high ambiguity, several plausible causes, and a real cost if the organization settles on the wrong one and moves on.

## THE DESIGN QUESTION, NOT A RULE

The useful move isn't a blanket policy in either direction. It's a specific question to ask before reaching for AI on anything genuinely ambiguous or consequential: does this decision need an independent first judgment before AI's view enters the room?

Sometimes the honest answer is no — the problem is routine enough, or low-stakes enough, that AI-first is simply efficient, and treating it otherwise would be its own kind of poor judgment. But sometimes the answer is yes, and that's worth deciding on purpose, before opening the tool, rather than by default, because the tool happened to be the fastest way to get any answer into the room. Once AI's explanation is already sitting there, coherent and plausible, an independent view becomes harder to produce. You can deliberately create alternatives, bring in a second reviewer who hasn't seen the first answer, or ask for competing hypotheses. But if genuine independence matters, the simplest design is to preserve it before the first answer enters the room.

AI doesn't only influence the answer. Sometimes it influences the thinking that happens before you reach one.

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**