

ESSAY 03

BETTER JUDGMENT IN THE AGE OF AI

AI Can Shorten the Time to Decide. It Cannot Shorten the Time to **Know**.

"AI can shorten the time required to decide. It cannot shorten the time required to know."

AI Can Shorten the Time to Decide. It Cannot Shorten the Time to Know.

By Alen Mayer

AI has gotten remarkably good at making decisions faster.

Research that took three days now takes minutes. Analysis that took a week happens before lunch. Comparisons across dozens of options, coordination across departments that used to require six meetings — increasingly handled in the time it takes to ask.

That's not a small thing. Most organizational delay has always been processing delay: gathering the facts, weighing the options, getting the right people in a room to agree. AI is genuinely dissolving that kind of waiting, and there's no case for nostalgia about it.

But there's a second kind of waiting AI hasn't touched at all, and conflating the two is where the trouble starts.

■ TWO KINDS OF WAITING

Call the first kind processing time — the waiting caused by how long it takes people to gather, compare, calculate and coordinate. AI collapses this relentlessly, and every quarter it gets better at it.

The second kind is different in a way that has nothing to do with how fast anyone works. Sometimes you're not waiting for an answer to be calculated. You're waiting for reality to produce evidence that doesn't exist yet. A material hasn't been tested under stress. A customer hasn't lived with a product long enough to reveal whether they'll stay. An organization hasn't had time to show whether people actually behave differently under a new structure, or just say they will in the meeting where it's announced.

There's a difference between waiting because people are slow and waiting because reality hasn't answered yet.

AI can shorten the first. It can also help us make better predictions while we're waiting for the second — using prior data, simulations, leading indicators, analogous cases. But it cannot turn an outcome that hasn't happened yet into an observed outcome. The evidence reality will eventually produce still takes its own time to exist.

■ THREE PLACES WHERE THE EVIDENCE HASN'T ARRIVED

A production line needs a replacement material, and an AI system can identify a qualified alternative and authorize the switch in minutes. But the required stress test takes twenty-four hours, because the material has to remain under defined conditions for that period before the result exists. No amount of faster authorization changes what that result will be. The decision got faster. The material's behavior did not.

A product team can ask an AI system to estimate the likelihood that a new feature increases retention, and get a confident, well-reasoned answer in an afternoon — probably a genuinely useful one. But an estimate of what customers will do is not the same thing as watching what customers actually do once they've lived with the change for a few billing cycles. The forecast can be excellent and still be a forecast. The gap between a strong prediction and an observed outcome doesn't close just because the prediction arrived quickly.

A leadership team restructures how two departments work together, and within days an AI system can surface every early signal available — sentiment in internal messages, meeting frequency, who's looping in whom. All of it real, all of it immediate. None of it is the same as knowing how people actually behave once the announcement stops being new and the change becomes just how things work now. That kind of evidence has a shape that can't be rushed, because it's made of people adjusting, testing the edges, and settling into habits — a process that runs on its own clock regardless of how fast the analysis around it moves.

Three different kinds of evidence — physical, behavioral, organizational. Each one takes its own kind of time to exist, and none of that time is processing time.

■ THE TRAP AI'S OWN SUCCESS CREATES

Here is the uncomfortable part. The better AI gets at collapsing processing time, the more naturally people start treating *all* delay as processing time — because everything around the delay got faster, so the one thing that didn't starts to look like the problem.

When research took three days and analysis took a week and coordination took six meetings, a twenty-four-hour test or a full product cycle before real data arrived didn't stand out. It was one slow thing among many. Now that everything else happens in minutes, that same twenty-four hours, that same product cycle, looks like the bottleneck. Someone eventually asks why this still takes so long.

Sometimes that's exactly the right question, and the honest answer is bureaucracy — a review nobody actually needs, a sign-off that changes nothing, a habit nobody's questioned in years. Strip it out.

But sometimes the honest answer is that this is where the evidence comes from, and there was never a faster way to get it. AI's success at removing the first kind of waiting

doesn't tell you which kind of waiting you're looking at. It just makes the remaining kind more visible, and visibility gets mistaken for inefficiency more often than it should.

There's a difference between waiting because people are slow and waiting because *reality* hasn't answered yet.

■ WHAT TO DO WHILE YOU WAIT

The answer isn't simply to wait. That treats every case as identical and makes patience into a policy instead of a judgment. It also isn't to override the wait every time speed is available, which is exactly the mistake the trap sets up.

The more useful set of questions looks like this: What evidence doesn't exist yet, specifically? What can you learn without all of it — a smaller test, a partial rollout, an early signal that's genuinely informative even if it isn't final? What can you commit to right now regardless? And if you act now, how reversible does that commitment stay while the missing evidence is still on its way?

Reversibility is one of the most useful levers here, not the only one. Some decisions can be staged, piloted, or held open in ways that let you move now without foreclosing what the evidence might later tell you. But some decisions genuinely can't be unwound once made, and knowing that in advance is itself valuable information — it tells you the twenty-four hours, or the product cycle, or the settling-in period, might be worth the wait after all.

Sometimes the right move really is to wait. The point was never to avoid waiting. It's knowing precisely why you're doing it — friction to be removed, or evidence still being made.

The faster AI becomes, the more important it becomes to know which kind of waiting you're looking at. Some of it is friction. Remove it without hesitation. Some of it is evidence, still being created by reality on its own schedule. Decide what you can before it arrives, keep what you can reversible, and know exactly what you're risking if you don't.

AI can shorten the time required to decide. It cannot shorten the time required to know.

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