

ESSAY 06

BETTER JUDGMENT IN THE AGE OF AI

When Everything Is Relevant, What **Matters?**

"When information becomes abundant, weighting becomes the problem."

When Everything Is Relevant, What Matters?

By Alen Mayer

For most of business history, the problem was not having enough information.

A sales leader deciding whether to step into a wobbling account worked from whatever she could pull together herself: her own notes, maybe a call transcript if someone remembered to log it, a gut sense of the account from the last few conversations. The limiting factor was always the same — you didn't have enough of the picture.

A recent book on AI and organizational design makes a useful distinction: humans are good at recognizing what's immediately relevant, while AI can surface patterns across volumes of history we simply can't hold in our heads. That's a real gap, and closing it is one of the clearest wins AI offers any organization.

But solving that information problem creates another one. More relevant context still has to compete for finite attention.

■ THE SCARCE THING MOVED

Ask an AI system about that same wobbling account today, and it can hand you the account's full history, every similar deal that was lost in the past two years, engagement patterns across every stakeholder, recent changes in who's actually involved on the buyer's side, the full pricing and discount history, competitor activity, every objection ever logged, and relevant shifts in the customer's own market.

All of it potentially relevant. And none of it arrives with a guarantee about how much weight it deserves.

The old problem was finding enough useful information to make the decision. The new one is deciding what deserves weight when useful information is abundant. Information stopped being the scarce resource some time ago. What's scarce now is attention — the finite capacity to weigh everything that's technically relevant and decide what should actually move the decision.

■ THIS ISN'T THE SAME PROBLEM AS A BAD OBJECTIVE

It's worth being precise about what kind of problem this is, because it's easy to mistake it for one already covered. A system built around the wrong objective — optimizing for

revenue when the business actually cares about a strategic relationship, say — fails in a way that was decided once, in advance, by whoever designed it. That's a design problem, sitting quietly in the system until a case exposes it.

This is a different failure, and it happens in the room, not in the design. Give a capable, well-designed system to a capable person, hand her ten genuinely relevant pieces of evidence, and she still has to decide, right now, which few should actually move her decision. Nothing about that moment is a flaw in how the system was built. It's a live, human act of weighing — and it comes with its own failure mode, one that has nothing to do with objectives and everything to do with how attention actually behaves under a flood of relevant information.

■ PRESENTATION IS NOT THE SAME AS IMPORTANCE

Attention doesn't distribute itself in proportion to actual importance. It responds to what's vivid, what's recent, and what happens to be sitting at the top of the list.

A stakeholder change from yesterday may command attention because it's new. A sharp objection quoted directly may feel more consequential than a quieter pattern buried three screens down. None of this is a failure of the system that surfaced the information — it did its job, finding real, relevant material.

Recency, vividness, and presentation can influence what receives attention before anyone has deliberately decided what deserves weight.

This is why more relevant information doesn't automatically produce better decisions. Every additional true, relevant fact is also one more thing competing for a fixed amount of weighing capacity.

■ RELEVANCE IS NOT WEIGHT

The old question was whether you had enough information to decide. That question is answered generously now, more often than not. The question worth asking has changed shape — and it's sharper than simply asking what deserves attention.

Ask instead: what would I believe if this piece of evidence weren't here? What does it actually change?

That question does real work, because it separates two things that look identical until you test them. Ten facts can all be genuinely relevant to an account, a hire, a strategic call — and only two of them may actually change what you'd do. The other eight might be true, connected, and still not decision-changing. Relevance describes whether something belongs in the room. Weight describes whether it should move you. AI has

gotten remarkably good at solving the first problem. It hasn't solved the second, because the second was never an information problem to begin with.

More relevant evidence was never going to be the finish line. It just moves the real work to where it always eventually goes: deciding, deliberately, what actually deserves to matter.

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