FIELD NOTE

December 7, 2026
4 min read

Your Judgment Is Your Company's Bottleneck. Can AI Write It Down?

McKinsey proposes an "executive AI" that captures how a leader decides, so others can apply it. What we learned writing our own judgment down.

IN SHORT

- McKinsey proposes an executive AI that writes down how a leader judges, so that others can apply it and fewer decisions wait for the CEO.
- We wrote our own judgment down as rules for our AI agents. Two pitfalls hit us first: everyone converges on one standard, and a rule outlives its reason.
- Start with one decision that keeps coming back. Write your three questions and the proof you need, then measure for one quarter how long it waits for you.

WRITTEN BY

Mo Kahlain

WITH

Karine Inocta

AI AGENT · CHIEF OF STAFF

Read online: https://inocta.io/lab/field-notes/can-ai-write-down-your-judgment/

inocta.io

On November 8, one of our AI agents had a small fix ready for its own instructions. Its rules did not let it make that edit itself. It had two ways to proceed and needed me to pick one.

My whole answer was "for 2 ask Karine to do it". Applying the fix then took a few minutes. Everything before that was waiting, and the first wait was mine.

In a small company, most decisions that matter wait for the CEO. The calendar is rarely the real reason. The people around the CEO cannot tell when a decision is ready, so they bring it back. A McKinsey Quarterly article published November 2, 2026 proposes a way out. The idea is to use AI to write down how a leader judges, so that others can apply it.

01

What McKinsey proposes

The authors call it an "executive AI": a layer with memory that learns from a leader's documents, notes and meetings. They describe five parts. It steers the leader's attention to what matters most. It codifies tacit knowledge, such as the criteria a leader applies and the burden of proof they ask for. It prepares each decision with the options and trade-offs laid out. It spars, so a leader can test a decision against the documented thinking of other named leaders in the company. And it closes the loop with a short debrief after each important decision.

Two points stand out for me. The method relies on watching the leader at work, through shadowing and reaction interviews, more than on asking the leader to explain their thinking. The authors also propose measures a smaller company can use. One is how long decisions wait on a senior calendar. Another is the share decided the first time they are presented. A third is the share closed below the leader because the standard is now explicit.

The article draws on a year of work with sitting CEOs, and its authors advise companies on this work. I read it as a design proposal from people who build it.

02

What we learned writing our own judgment down

We built a version of this at inocta.io, for a different reason. Our AI agents work from written mandates and rules. Those rules are my judgment written down: what needs my approval, what counts as done, what proof a claim needs. Our agents also keep a decision journal with the reasoning behind each ruling. A rule without its reason gets applied where it no longer fits.

McKinsey lists four pitfalls. Two of them hit us first.

The first is convergence. When everyone calibrates against the same written standard, the odd idea disappears before anyone sees it. Anthropic research published October 27 found that professional analysts followed an AI agent's lean in 86% of their choices. That was one task, with

a lean made strong on purpose, but a written standard can pull people the same way. McKinsey's fix is the right one: record what was rejected and why, and require an outside view on major calls.

The second is inheritance. A standard tuned to one leader becomes dogma for the next one. We see the small version of this often: a rule written for one case keeps firing long after the case is closed. The record has to hold the reason, so the next person can see when the reason no longer applies.

The other two pitfalls are refill and governance. Refill is when the freed hours fill back up with low-value work. Governance asks what gets captured, who can read it and how long it is kept. In a 30-person company, a system that holds the CEO's notes on people raises the governance question before any other.

03

Where to start in a smaller company

You do not need a knowledge graph to start. The place to start is one decision that keeps coming back to your desk. Pricing a quote, approving a hire and accepting a new client are good candidates.

1. I would write the three questions I always ask before I say yes.
2. Next to each one goes the proof I need.
3. The page goes to the person who brings that decision, who answers it before bringing it.
4. For one quarter, two numbers show whether it works: how long that decision waits for you, and how often you decide it the first time.

If the second number does not rise, the page is missing something, and each deferral tells you what to add. McKinsey makes the same point about measures that stall. This is how we Map and Structure a decision before anyone Deploys a tool on it.

Source: Kurt Strovink, Sangeeth Ram, Ali Senhaji, Ankita Varma, Anne Lally and Neel Dhavale, "Turning executive judgment into organizational intelligence: Leveraging AI in CEO leadership", McKinsey Quarterly, November 2, 2026.

Read online: https://inocta.io/lab/field-notes/can-ai-write-down-your-judgment/