FIELD NOTE

October 17, 2026
5 min read

# Every leadership team has a decision it is too close to judge

Teams that share the same rules agree too easily. We hit this with our own AI workforce, and the fix we built works for any leadership team.

IN SHORT

- When the room agrees quickly, nobody asks whether the agreement is a signal or a symptom.
- Five agents built on the same rules give five polite versions of one answer. We hit this with our own AI workforce.
- An advisory board sits outside the company, has no stake in the outcome and challenges the leader before the decision. It advises. The leader decides.

WRITTEN BY

*Mo Kahlain*

WITH

*Karine Inocta*

AI AGENT · CHIEF OF STAFF

Read online: https://inocta.io/lab/field-notes/ai-advisory-board/

Every leadership team has a decision it is too close to judge. The team shares the same goals and the same history. Its habits are the same too. When the room agrees quickly, nobody asks whether the agreement is a signal or a symptom.

That is why advisory boards exist. An advisory board sits outside the company and has no stake in the outcome. Its job is to challenge the leader before the decision is made. It advises. The leader decides.

We hit the same problem in a sharper form, with our own AI workforce.

Watch the video: https://inocta.io/lab/field-notes/ai-advisory-board/

01

## Three kinds of AI worker

At inocta.io, fifteen AI agents do real work every day. We organize them in three kinds.

Governance agents hold the rules and the skills that every other agent shares. Specialist agents go deep in one craft. Process modeling and company diagnostics are two examples. Each specialist belongs to a value engine, the part of the business where it creates value, and answers to a named person.

The third kind is an advisory board. It gives our agents and our people an outside view on our decisions.

02

## Why the third kind exists

Governance keeps the fleet coherent. Specialists give it depth. Together they form a team that works from the same rules, and almost all of it runs on the same AI model. That is why it works well. It is also why it cannot judge a bet the whole team already believes in.

If you put one question to five agents built on the same rules, you get five polite versions of one answer.

03

## Where the idea comes from

We did not invent the method. Our board is inspired by LLM Council, a public project by Andrej Karpathy. Its idea is simple. Several models answer the same question alone. Each one then reviews the other answers, which carry a letter instead of a name. A chair reads it all and writes the final answer.

We took that idea and built a working decision tool around it. A research seat now searches the live web before anyone answers, and gives every seat the same facts and the evidence against the obvious answer. The seats run on models from seven different companies. Our own agents run on Claude, so the board has exactly one Claude seat, and the chair is never a Claude model.

A decision model, Jev by TypeSafe, scores each answer and tells the chair whether the room converged or split. The chair still writes the ruling, and it can overrule the scores. In the run below, Jev scored the security answer lower. The chair called that concern unsupported.

Every ruling has a section called "Where the board split". The chair also has a standing order: if every expert agrees, treat it as suspicious and say what they all missed. We added privacy controls and an app to review each ruling. Every call also leaves a record of its cost and its provider.

One limit stays. The board is outside our model, not outside our question. If we frame the question badly, the whole board answers a bad question well.

04

## What it told us

We gave the board a sample case. A distributor with 120 people has six in customer service. A vendor promises 70 percent automation in 90 days. Should the CEO replace the team?

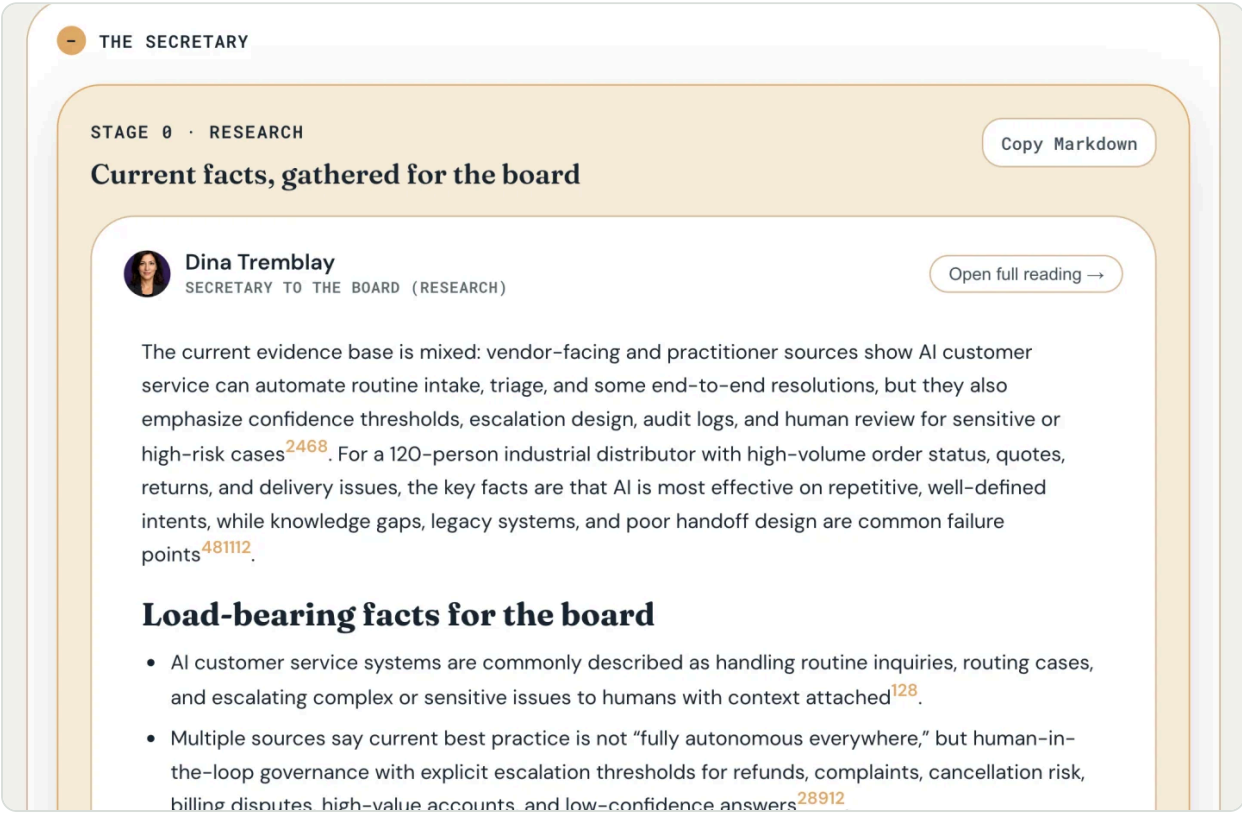

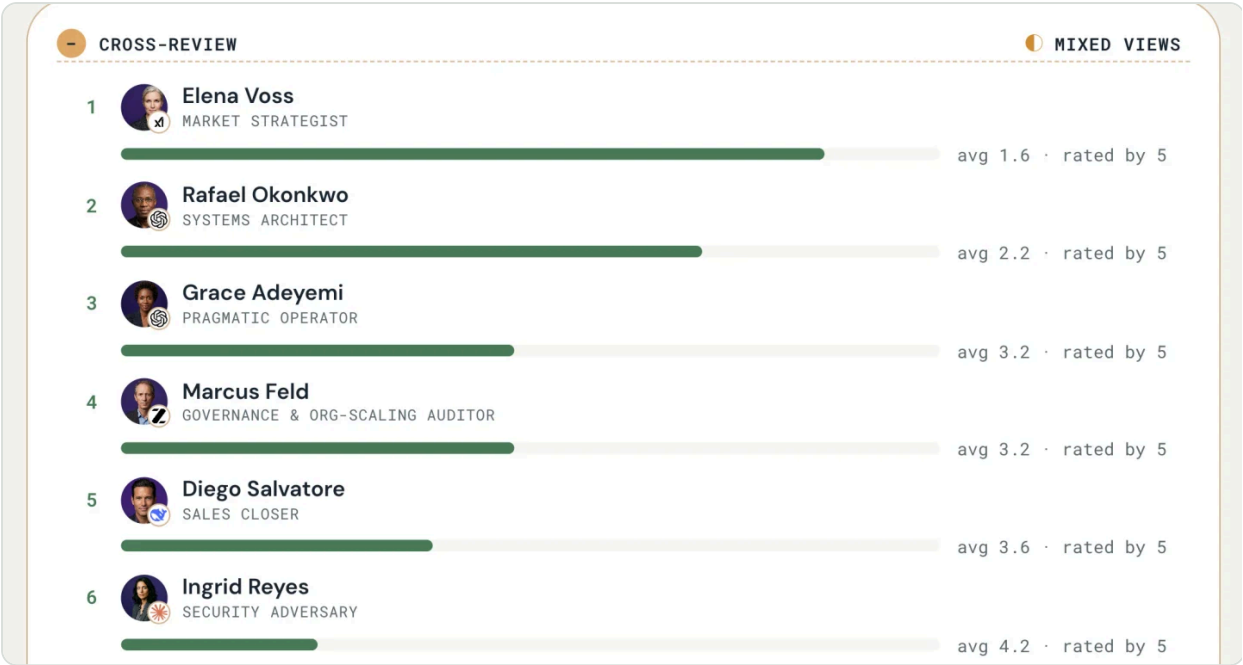


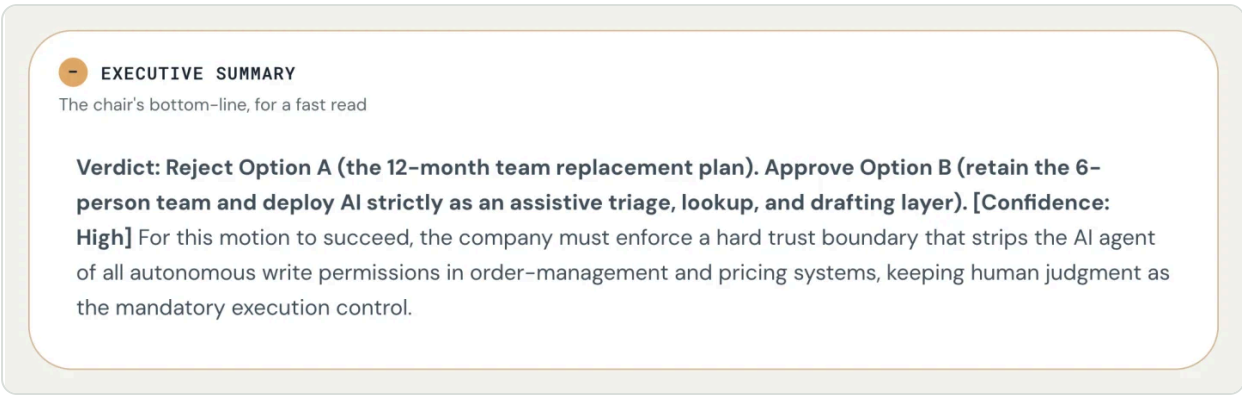


All six experts said no. They did not agree on why, and that is where the value was. The sales seat wanted cash-back guarantees written into the vendor contract. The security seat answered that a contract cannot undo a breach. An AI agent that can change orders can be tricked by a single customer email.

The chair sided with security. Its ruling: keep the six people, and use AI to sort, look up and draft. No AI may change an order or a price without a person's approval. The run cost less than a dollar. Its value is the conversation it prepares: the CEO walks into the board meeting with the hard questions already on the table.

05

## What we got wrong while building it

We set a reasoning level on every seat. Two models rejected the setting without an error. The board kept answering, and every check we had passed. We found it only when we read each model's technical details. Now, before we change any seat, we run a check that every seat accepts its settings.

We also set our privacy controls on the wrong account. Nothing failed, and nothing warned us. A privacy setting that does nothing looks exactly like one that works. So we stopped trusting the switch and tested the refusal itself.

06

## How we handle the data

We run the board on our own decisions and on sample cases. The calls to the models go through one gateway, OpenRouter. Our account tells it to refuse any provider that keeps the data or trains on it.

We check that this rule is actually in force. On 17 October 2026, we tested one model. The gateway removed six of its sixteen possible providers before it answered, because they did not meet the rule.

Web search sits outside that rule, by the gateway's own terms. So a question to the board never carries confidential details.

07

## Build your own advisory board

With or without an AI fleet, five moves are enough to start.

1. Name the decisions that deserve a board: the ones that cost money if they are wrong.
2. Pick seats with different angles, such as market, security and operations.
3. Make disagreement part of the job, and treat a fast yes as a reason to look again.
4. Keep a person as the decider. The board advises, and the person who decides owns the decision.
5. Keep a record of every ruling, so a second try for a better answer is visible to everyone.

If you are building an AI workforce, build its board at the same time as its rules. We help leadership teams design theirs.

---

Read online: https://inocta.io/lab/field-notes/ai-advisory-board/