

The Five Laws Series

Communication #7

Ten Copilots, Ten Realities

 **Causum**

A steering committee expects to decide the future of a critical vendor—until two AI-generated reports from different departments confidently describe opposite realities. As the meeting unravels, executives discover a deeper threat: their teams no longer disagree about decisions, but about reality itself. In a world of personalized AI, can an organization still share a single truth?



Nobody chose this

No one decided to fragment the company's picture of reality. It arrived as a side effect of features.

Personalization tailors each person's results. Summarization compresses differently for different prompts.

Retrieval surfaces different documents for different phrasings of the same question. Each feature is individually useful; their sum is that each employee now works inside a slightly different rendering of the company, and the renderings never reconcile — because *reconciliation was never anyone's feature*.

Institutions run on shared premises. Strategy, escalation, audit, post-mortems: all of them assume that two competent people looking at the same facts see approximately the same situation. AI mediation quietly dissolves that assumption while every individual tool performs to spec.¹


Views may differ. Facts may not.

How sensitive are AI verdicts to nothing more than phrasing? A controlled study measured exactly this: regulated decision situations, each retold five ways — same facts, same right answer, different wording — through the same AI system.²

The ungoverned system changed its verdict on 9 of the 61 situations based on wording alone. One in seven. Same facts, different phrasing, different answer. Now recall that your teams never even use the same wording: different people, prompts, and contexts guarantee it. The committee-meeting collision isn't bad luck. At one-in-seven sensitivity, it's arithmetic.

The same study measured a governed alternative on identical retellings: *61 of 61 identical verdicts*. Not because that system was smarter, but because of *where its answers came from*: each verdict was derived from a typed record of the facts checked against a declared rulebook, and the derivation provably doesn't vary with phrasing. Wording can't move a verdict that doesn't rest on wording.

One more result closes the loop: the study also tested whether the fix was just neater inputs. Rendering the same modeled content in different styles changed nothing; what carried the improvement was modeling the domain — agreeing on what the facts *are* — and what finished the job was verification.² Shared structure, not shared phrasing, is what aligns judgment.



**THE MEETING
WASN'T DECIDING
WHAT TO DO.**

**IT WAS DECIDING
WHAT WAS TRUE.**

Our teams don't disagree; their AIs do.

Let every team keep its own lens, summaries, and emphasis. But consequential determinations — is this vendor compliant, does this claim qualify, is this exposure material — should resolve against one governed record, where the same situation gets the same verdict no matter who asks, how they phrase it, or which tool they're holding. One spine, many windows. The spine is what the committee meeting was missing: a layer of determinations that both decks would have had to agree on before the hour started.

If you've walked out of a meeting like that one thinking *we didn't disagree about the decision, we disagreed about reality* — that's this episode, and you diagnosed it before your tooling did.

Sources

1. Kanjani AI Research (2026). *Declaration of the Five Laws of Mediated Intelligence Systems* (on divergence). AIGP v4.0.
2. Fatahi, R. (2026). *Natural Language Has No Formal Semantics. Verify the Decision Anyway*. doi:10.5281/zenodo.21303485

