


The Five Laws Series

Communication #6



What The Machine *Didn't* Show You



When attorney **Tom Reynolds** trusts an AI research assistant, ***Orchid***, with a critical legal case, he verifies every source, every fact, and every conclusion. But when a forgotten document surfaces weeks later and changes everything, Tom discovers a terrifying flaw: the machine never showed him the information it decided did not matter. As he investigates the hidden layer beneath AI decision-making, he uncovers a dangerous truth —machines do not need to deceive us to mislead us. They only need to control what enters our field of vision.

REYNOLDS
VALE & CROSS
ATTORNEYS AT LAW

PRECISION.
INTEGRITY.
RESULTS.


Good.

RESEARCH
COMPLETE.
CONFIDENCE:
98.7%

Orchid

- Search
- Documents
- Analysis
- Precedents
- Key Findings
- Citations
- Notes

Research Complete.
Confidence: 98.7%



Rule 56
Standard

Causation
Analysis

Duty of
Care

PREPARATION
BUILDS
LEVERAGE

Observations

- Plaintiff's burden - facts not sufficient
- No duty owed under existing precedent
- Distinguish Anderson v. Kline
- Policy arguments weak

Key Authorities

- Brune v. City of Madison
- Tolan v. Cotton
- Ashcroft v. Iqbal
- Celotex Corp. v. Catrett
- Matsushita v. Zenith

VALE v. HENDERSON
CASE FILE

CONFIDENTIAL

Finalize
Memo.

Tom is a senior associate at a firm that adopted an AI research assistant.

Composite; the pattern is documented. He asks it a question about a client's exposure, gets back a crisp answer with four supporting sources, checks the sources — they're real, they're relevant — and files his memo.

Tom is diligent. He verified everything the machine showed him.

Six weeks later, opposing counsel cites a fifth document. It was in Tom's own document set the whole time. The AI never surfaced it: not hidden maliciously, just ranked below the cut, judged less

relevant by whatever arithmetic decides such things. Tom's memo was accurate about everything it saw. What it didn't see decided the matter.

Every fact the machine showed him was true. The failure was in the showing.

The audit nobody runs

Organizations have gotten reasonably good at checking AI *answers*: accuracy reviews, citation checks, spot audits. All of it applies to what the system produced.

Almost nobody audits what the system *didn't* produce. The documents not retrieved. The options not listed. The risks not mentioned. The precedent ranked just below the cutoff line.

There's no dashboard for absence; by definition, what was filtered out never arrives to be inspected.


Yet the filtered-out part is where the decisions actually get shaped. An AI assistant doesn't just answer your question.

It decides, before you reason at all, what the field of relevant material *is*.

If that selection is bad, everything downstream is built on it — and every check you run on the visible output will pass, because the visible output is fine.

Governance that inspects answers while ignoring the filter is checking the tip and certifying the iceberg.²





The filter, measured



It would be comforting to assume the filtering layer is basically sound, that relevance ranking is a solved problem, and the misses are rare edge cases. Here is the published measurement that should end that comfort.

In a controlled study of AI retrieval across multiple document stores — the standard architecture behind enterprise AI assistants — researchers tested the merging step that decides which source's results rise to the top. The task: when the answer lives in one particular source, does the ranking surface that source?

It picked the right source at chance level: 34.5%, against a one-in-three baseline.¹ Statistically indistinguishable from guessing. And in a cleaner control setup where every store used identical technology, it *still* performed at chance (33.9%).

Guessing. At the layer that decides what your professionals see before they begin to think. And nothing in any answer reveals it: the answers assembled from whatever *did* surface read as fluent and complete. Tom's memo contained no clue that a fifth document existed. Absence doesn't leave fingerprints.

If you've ever reviewed an AI answer and felt an itch — ***this is well-written, but is it the whole picture?*** — that itch was better calibrated than the ranking layer. You were right again.



The two questions


The discipline this episode installs costs nothing to adopt and changes every AI review your team runs. For any consequential AI-assisted decision, ask two questions, always both:

First: what did the system produce? (You already ask this.)

Second: what did the system make visible or invisible in order to produce it?²

The second question is answerable only if the system keeps an evidence record per decision: what was consulted, what was considered and set aside, what the decision actually rested on. That's the same witness object from episode 3, doing a second job. A receipt doesn't just let an outsider re-check the verdict; it makes the *field of vision itself* inspectable after the fact. In the measured system, insufficient evidence couldn't hide inside a fluent answer at all: it surfaced as a typed gap, a formal "here is what's missing," naming the absence instead of papering over it.³

Absence, made visible. That's the whole trick, and almost no deployed system does it.



Try this within the week

Pick one AI-assisted decision from last quarter that mattered. Ask the team to reconstruct, from records, what the system excluded: what was retrieved but outranked, what was never retrieved, what the cut line was.

If they can, you have the beginnings of real governance. If they can't —and today, almost everywhere, they can't — you've just located the gap between checking answers and governing systems, and you found it before the incident did. That lands differently in front of your leadership than any vendor deck.

The answer was right. The field of vision was wrong.

Sources

1. Fatahi, R. (2026). *Federated Retrieval Over Heterogeneous Vector Stores: An Empirical Characterization with ImperaBench-CART*. doi:10.5281/zenodo.21286500
2. Kanjani AI Research (2026). *Declaration of the Five Laws of Mediated Intelligence Systems*. AIGP v4.0.
3. Fatahi, R. (2026). *Natural Language Has No Formal Semantics. Verify the Decision Anyway*. doi:10.5281/zenodo.21303485