

We Don't Know What We Don't Know. Neither Does AI

By Pascal HERNALSTEEN and Margaux HERNALSTEEN-LECLERCQ*

As a board member of a Luxembourg fund, you are about to approve the investment memo. The AI read 3,000 pages across 40 categories. The deal team is happy. Risks have been flagged, the recommendation is soundly argued, and the financial tables look clean and cross-referenced. Someone on the team says the words that should worry you: "It looks like nothing was missed."



The GP has built AI into the deal workflow because it works. It reads faster, it indexes every claim back to a source. But the memo tells you what the AI read. It says nothing about what nobody put in front of it. And that gap has no direction you can predict. That's the part that should keep you up before signing. Let us introduce three characters. The most famous is easy to kill.

The first character: hallucination

Everyone knows this one. The AI invents a fact that reads as true. A contract that doesn't exist, a growth rate nobody computed, a compliance status the target never held. It's the loud problem, and it's the easy one, because it's mechanical. You build three routines into your review protocol and it mostly goes away.

Require a cited source for every factual claim. Spot-check a dozen of the critical numbers against the underlying documents. And where a figure is built from several inputs, ask for the inputs and the calculation, not just the answer. Even better, build an agent that will review all numbers and claims against the source. Do that, and hallucination stops being your problem. So we spend most of our time on the character that gets almost no airtime.

The second character: silence

The AI works on what you feed it. It can't miss what was never uploaded. Think about what a data room doesn't contain.

- The customer contract that was "available on request" and never requested.

- Eighteen months of CSSF correspondence that simply isn't there.
- The environmental appendices missing from an industrial asset.
- The suspicious-activity report that have gone to the Financial Intelligence Unit and left no trace in the file.

The AI won't flag any of these. It had nothing to read. You can't summarise an absent document, and you can't price a risk that never reached the room. A person who has run twenty of these deals notices the empty slot. They walk in with a mental model of what a complete data room looks like, so the missing folder stands out. The machine has no such model. It's graded on the exam it was handed, and nobody checked whether pages were torn out of the exam. This is the dog that didn't bark. In Conan Doyle's "Silver Blaze," Holmes cracks the case on one fact: the guard dog stayed silent during the theft. It knew the thief.

The absent bark was the clue. In diligence, your most expensive miss is usually the document nobody requested and the question nobody put to the seller. The fix is sequence. Put a human in front of the machine. Have an industry expert audit the folder structure against a standard deal checklist first, before a single prompt runs. The expert marks what's absent. Then the AI analyses what's present. Once the memo looks complete, nobody goes looking for the hole.

The third character: the traceability trap

This one is the sneakiest, because it wears the badge of the AI's biggest strength. Traceability is real and it's



powerful. Every assertion in the memo links back to a section in a source document. You can audit any claim in seconds, but that trace is only as good as the version it points to. Here's a case we saw. The CFO sent a trial balance showing EBITDA of 33.1 million. Two days later, he emailed a correction: typo, the real figure is 31.3 million, revised document attached. The AI built the memo off the first file. Trace the number in the data room and it leads you straight back to the original, the one with the typo, with total confidence and a clean citation. The trace was perfect. It just pointed at the wrong truth.

Two disciplines close this. First, feed the AI everything, not just the CIM and the data room, but every Q&A session, every email exchange, every correction memo. The version you analyse has to be the latest one. Second, ask the AI to recompute every calculated figure from its components and flag any discrepancy across documents. The gap only shows when you read across files, which is exactly the work a human skips and a machine can do well, if you point it there.

Why all three are the same problem

Here's the thread that ties them together, and it's the reason none of this is a solved problem. Each of these three can push your valuation in any direction. Hallucination can invent a strength or invent a clean bill of health. Silence can hide a liability or bury real upside nobody surfaced. A stale trace can lock in a number that's too high or too low.

So the error has no reliable sign. You're wrong in a direction you can't predict. And a bias you can't see is one you can't correct for. A known bias, you adjust. An unknown one, you sign. The market debates whether AI makes deal teams too bullish or too careful. The unsigned error matters more, and it's why the human seat next to the machine hasn't gone anywhere.

The credibility problem underneath

There's a reason this slips through even sharp teams. An analyst memo is easy to challenge. You expect a junior to have gaps, so you ask for the backup and the reconciliations. An AI memo wears a better suit. Cited sources, structured risk sections, clean tables, every angle apparently covered. That polish is exactly what disarms the room. It looks pre-verified, so the

committee nods. The confidence is genuine. The completeness might not be.

1. Fiduciary duties

Directors owe the fund care and judgment, and they owe it personally. The reading can go to a machine. The duty stays in the boardroom. When the regulator or the LPs come asking, "the AI reviewed it" will not be an answer.

2. The board's role hasn't changed

The technology now drafts the whole recommendation. Rationale, risk section, conclusion. Fine. The deal team's job, and the board's after it, is the same as it was in 1995. Challenge, and ask the question the memo skipped.

3. Critical thinking in the age of AI

Critical thinking is the scarcest resource in the room. Three questions we'd put on the table before any IC. What did you expect to find in this data room that wasn't there? Who checked the folder against a real deal checklist, and what came up missing? And did you prompt the AI to attack the thesis, or only to build the case for buying?

That last one matters most. The best committees we've seen use the machine to generate doubt, not comfort. Build the bear case. Tell me why we walk. Which numbers fight the seller's story. What would a hostile IC member hit first. Why is this business in trouble in three years. That's where the tool earns its seat. As a generator of the questions a tired human stops asking at page 400.

When the portfolio company wobbles in year three, the board will ask whether the right document was in the room, and if it wasn't, whether anyone thought to ask.

The dog that didn't bark still won't. Someone in the room has to notice the silence.

** Pascal Hernalsteen is a Luxembourg-based operating partner and independent consultant specialising in fund services M&A, post-acquisition integration, and AIFM governance.*

Margaux Hernalsteen-Leclercq studies Data Sciences and Business Analytics at ESSEC Business School and CentraleSupélec in Paris, where her research focuses on the intersection of AI tools and financial decision-making.