Interpretive Drift

How AI Speeds Up a Leadership Failure

In late 2020, Peloton's CEO was asked point-blank on an earnings call how the company was managing the risk that pandemic-driven demand might normalize. His answer was unambiguous: "When you say 'normalized coming out of covid,' we don't see that."

Eighteen months later, the stock had lost more than 90% of its value.

This wasn't negligence. Peloton was responding rationally to a logic its own results had confirmed, quarter after quarter. That's what made the collapse so fast and so hard to reverse: demand normalized more or less exactly as the signals suggested it would — but by then, the organization could no longer tell what those signals meant. The data was there. The analytical horsepower was there. What failed was the interpretive system that turned data into meaning.

Nokia saw smartphones coming. Kodak invented digital imaging years before the market caught up. In every case, the information was present. The interpretation was foreclosed — not by failure, but by success.

This slow erosion of alignment between what the leader believes creates value and what actually creates value now — concealed by the very results that seem to prove the old belief still holds—is what I call leadership interpretive drift.

And AI is making it worse, faster, on two fronts most leaders only see one of.

Why success is the problem

Here's the mechanism, and it's not about attention or effort.

When a strategy gets validated by results, it stops behaving like an assumption you test. It becomes the lens you see through. Prahalad and Bettis had a name for this decades ago — "dominant logic" — and it's not a flaw in how organizations work. It's how they build the coordination and speed that let them execute at scale.

The problem is that the same process that builds that capability also quietly dismantles your capacity to question it. The better you get at running your current strategy, the worse you get at noticing when it stops working. Peak performance and peak interpretive rigidity arrive at the same moment, in the same organization, for the same reason.

You've felt this. A concern gets raised in a meeting and the room moves past it. An anomaly in the numbers gets explained away. A case for a different direction feels obvious to you and completely unmakable to everyone else. That's not a failure of attention. That's the architecture doing exactly what it was built to do — and it was never built to question itself.

The usual fixes don't touch this. Faster review cycles just reassess the same assumptions faster. More data sharpens conclusions inside the frame you already have. More flexibility adjusts how you execute, not what you're executing toward. None of it asks the one question that matters: is the underlying premise still true?

AI's two fronts

Most conversations about AI and strategy stop at the first front and never get to the second — which is the one that should worry you more.

Internally, AI reinforces your current logic faster, more thoroughly, and with more apparent objectivity than any human process ever could. A system trained mostly on your own organization's data will confirm your organization's own thinking — with a rigor that makes it harder to question, not easier.

Externally, AI is compressing how long any value logic stays valid at all, because your competitors and new entrants can now test and deploy alternative ways of creating value at a speed no earlier technology allowed. Interpretive drift has always meant your governing assumptions falling out of step with how value is actually being created around you. AI doesn't create that gap — it shrinks the runway you have to close it.

Legal services show both fronts running at once right now. AI isn't just making contract review faster inside the existing associate-leverage model — it's generating a different model altogether, one where expertise is necessary but no longer sufficient. And the firms most exposed aren't the weak ones. They're the most successful — the ones most confident their model still works.

What this asks of you as a leader is two things, working together. First: treat strategy as a provisional hypothesis, and question it while it's still producing acceptable results — not after the evidence has already closed off your options. Second: recognize that your own judgment won't survive contact with an organization built to reinforce the old logic. You can see the shift coming and still get overridden by your own company, unless you've built something that supports reinterpretation instead of suppressing it.

Judgment without structure gets overridden. Structure without a leader who'll use it reverts to default under pressure. You need both.

So what does that structure actually look like?

The 4D framework: building an organization that can question itself

Four dimensions, each targeting a different point where thinking quietly gets displaced by execution.

They all rest on one prior move: **separate your organization's identity from the strategy it currently runs.** When identity fuses with strategy, revising the strategy feels like self-destruction. When identity sits above strategy, you can change how you compete without destabilizing what you are.

Direction — Keep purpose stable, strategy provisional

The most common way drift becomes irreversible isn't dramatic. Strategy just gradually becomes identity. Once that happens, questioning it feels existentially threatening instead of useful, and the people who raise concerns get read as not understanding the business.

Google invented the Transformer architecture that powers today's large language models. It had the talent and the data. What it didn't have, by most outside accounts, was separation between what Google was and how Google currently competed — search wasn't just a strategy, it was identity and revenue model and the lens through which the AI threat itself got evaluated. When ChatGPT launched, reporting described a company-wide “code red,” with founders pulled back into strategy discussions they'd stepped away from years earlier.

Compare Microsoft under Satya Nadella, who's described the shift to cloud and AI not as an abandonment of what Microsoft had been, but as a deeper realization of it — a company that empowers people and organizations. Purpose absorbed the disruption. Strategy got revised. The organization could move.

Design — Build for reconfiguration, not just execution

Structures encode history. Over time they embed the logic of whatever worked before, routing resources toward established priorities and making it genuinely hard to pursue anything that doesn't fit that model — not because anyone's rejecting change, but because the structure literally can't accommodate it.

The fix isn't running more innovation initiatives. It's whether new understanding can become new action before the pull of “what already worked” drags it back into line — structures with boundaries that flex, resources that persist past the current business logic, and decision rights that let new approaches develop before they have to fight the old one for legitimacy.

Disciplines — Use execution to surface misalignment, not hide it

Strong execution is one of the best ways to conceal interpretive drift. When your assumptions are losing validity, disciplined delivery doesn't reveal that — it confirms the logic that's already degrading. Targets get met. Variance gets explained away. You execute well against a reality you no longer fully understand.

Jensen Huang's decade-long bet on CUDA at NVIDIA is this discipline in action. When he started funding it in 2006, NVIDIA was a gaming graphics company betting that parallel computing would become foundational to scientific and industrial workloads. Nothing confirmed that bet for six years — until 2012, when deep-learning researchers discovered GPUs were exactly the hardware needed to train neural networks at scale. NVIDIA's preparation became the infrastructure of the AI revolution.

Huang ran strategy like a scientist: a falsifiable hypothesis, investment sustained ahead of validation, resistance to chasing whatever the current numbers confirmed. When AI generates strategic analysis for you today, it draws on historical patterns and confirms established logic by default. Treat that confirmation as a reason to look harder for what disagrees with it — not as the answer.

Development — Preserve the capacity to see what doesn't fit

Even a well-built Direction, Design, and Disciplines depends on something you can't engineer directly: people who can recognize what doesn't fit and act on it under real uncertainty. That capacity erodes quietly, through ordinary promotion decisions. The leaders who deliver within the current logic advance. The ones who challenge it get managed out. You end up with exactly the leaders your current logic wants — and progressively fewer who could question it.

When Alan Mulally arrived at Ford in 2006, disagreeing openly with prevailing assumptions was career-limiting. His weekly Business Plan Review — status reported in green, yellow, or red — made hiding problems more costly than admitting them. The turning point came when Mark Fields reported a red on a major launch and Mulally applauded him for it instead of punishing him. Within weeks the room was full of yellows and reds — not because performance had gotten worse, but because the willingness to say so had come back. It hadn't been missing. It had just been too expensive to use.

As AI gets better at synthesizing complexity and modeling scenarios, the leader who's valuable mainly for analytical throughput is now competing with a system that does that faster. The leader valuable for judgment — recognizing when the model has stopped working, acting under real uncertainty — offers something no AI system replicates. That gets more valuable, not less, as AI absorbs more of what leadership used to require.

How the four fit together

Direction makes reinterpretation safe. Design lets it become action. Disciplines make misalignment visible instead of burying it. Development ensures you still have people who'll act on what surfaces. Pull any one out and the other three collapse under normal performance pressure. The difference between organizations that renew and organizations that drift isn't talent or resources — it's whether this was in place before the moment of maximum consequence arrived.

Six things to actually do

The framework is what your organization needs to have. These are what you, personally, need to keep doing to keep it functioning — especially in the moments when executing the current plan is most tempting and questioning it is most costly.

1. **Separate what you stand for from how you currently compete.** If your organization can't describe its purpose without describing its business model — “we win through manufacturing excellence” — strategy has already become identity, and it can't be questioned cleanly anymore. Say what you stand for independently of how you currently compete, and invoke that separation on purpose when strategy comes under pressure.

2. **Treat your structures as hypotheses, not infrastructure.** Structures feel permanent because they're expensive to change and current success seems to validate them. Ask, periodically, of every major structural arrangement: does this still serve where we're going, or mainly the logic we used to run on? Resource allocation that keeps favoring the existing business is an active bet the old model still holds.

3. **Run your strategy like a scientist, not a believer.** Name the assumptions your strategy depends on, explicitly. Design your execution system to surface evidence those assumptions are weakening — not just evidence that execution is strong. When AI-generated analysis reinforces what you already believe, treat that as a prompt to search harder for what disagrees with it.

4. **Make execution your early warning system.** Most performance reviews only ask: are we delivering against plan? That's half the question. Add a second one: what are we learning about whether our assumptions still hold? Where did strong execution produce weaker results than expected? What anomaly did we explain away that deserved more attention?

5. **Build the organization around productive disagreement.** As organizations succeed, they get less tolerant of dissenting interpretations — not from malice, just from the same selection pressure that produces high performance. Reversing it takes more than good intentions: designated challenge roles in major reviews, investment theses that name what would have to be wrong for the thesis to fail.

6. **Develop leaders for judgment, not just performance.** Ed Catmull's Braintrust at Pixar is the model here — a structural mechanism for challenging a film's creative premises as a normal part of production, with no authority to force changes and full accountability to say what was genuinely seen. The result was one of the great creative runs in film history, because the organization was built to question its own premises before the market forced it to.

The moment this asks of you

Governing premises don't fail suddenly. They erode while performance still looks fine — reinforced by the very success that makes them hard to question. The real leadership challenge isn't recognizing failure. It's acting while success is still giving you every reason not to.

If you're reading this, you probably already sense it somewhere at the edge of your attention — something in the logic behind your current results is shifting, and you're choosing, rationally, to wait for more evidence before you act. That wait has a cost. The longer you delay, the narrower your options get, and the more the eventual change looks like crisis response instead of a choice made with room to maneuver.

The hardest moment to revise a strategy is exactly the moment it's still working. As AI keeps compressing how long any strategy stays valid, that moment is increasingly the only one where you still have enough room left to act on it. If you want the full argument and the research behind it, the complete paper, *Interpretive Drift: Demanding a New Kind of Leader in a World That No Longer Stabilizes*, is on SSRN.