

The Harm Is In The Trajectory

*Every interaction was compliant. Every output was accurate.
Every action was authorised. And your people are not who they
were.*

CAUSUM COGNITIVE OFFLOADING · TEMPORAL-LINKING JUNE 2026

Your best analyst stopped checking
roughly eight months ago. *Nobody logged
it.*

No policy changed. No threshold was breached. No alert fired, because no instrument you own is watching for this. Every single one of those interactions passed every control you have — and each was genuinely fine, considered on its own, which is the only way your governance considers anything.

What changed was the person. Verification became a formality, then a habit, then an omission. Confidence detached from competence. Judgment migrated to the system so gradually that no one experienced a moment of handover.

They do not know it happened. Neither do you. And your logs will show, with complete accuracy, that nothing went wrong.

ABSTRACT

Governance is watching the wrong unit

AI governance is framed as control over isolated events: approvals, audits, thresholds, guardrails. This is a reduced practice — governance as a series of operational gates applied to discrete deployments.

The deeper problem is that human-machine interaction unfolds over time. People offload cognition to systems that shape how they decide, reason and act — subtly, cumulatively, and largely unseen.

Cognitive offloading is not pathology. It is a natural feature of human adaptation and tool use, and always has been. But **harmful dependency is undetectable without temporal-linking**: connecting moments across time, and practising second-order observation.

Governance must observe trajectories, not transactions.

THE CORE ARGUMENT

You secured the gates and ignored the weather


Governance did not shift from a hydraulic model to an atmospheric one. It has always contained both. Practice narrowed to one of them, and it narrowed toward whichever was easier to count.

THE HYDRAULIC DIMENSION What you measure Visible flows, channels, gates, thresholds, approvals, controls. It asks what is moving, where, who may move it, and which gate should open or close. Every one of these is trivially measurable. That is why it survived, and why it became the whole of the practice.	THE ATMOSPHERIC DIMENSION What you abandoned Uncertainty, pressure, incentives, culture, trust, fear, confidence, ambiguity, accumulated exposure. It asks what environment the human operates within, how it is changing, and how those changes shape judgment over time. None of it fits in a field. So it was dropped — not by decision, but by instrument.
--	--

The failure is not that controls exist. It is that in narrowing to the measurable, we stopped attending to the conditions and environment of a failure — which is where failures actually form.

Governance as praxis must reintegrate both: what shapes behaviour, the agents who navi-

gate, the long view of purpose, and the pathways our choices create over time.



REDUCED PRACTICE VERSUS GOVERNANCE AS PRAXIS. THE BUCKET IS FULLY INSTRUMENTED, FULLY COMPLIANT, AND GOING NOWHERE.

LOOK AGAIN AT THE LEFT PANEL

That is your governance programme

The bucket is not a caricature of bad governance. It is a picture of governance that is working exactly as designed — gates operating, pipework sound, console staffed, every reading nominal.

It contains a vessel that cannot go anywhere, watched by an operator who has confused the barrier for the voyage. The star is still on the drawing, but it has faded to an outline. Purpose did not get overruled; it stopped being an input.

Most people reading this believe they are the figure on the right. The instrument that would tell them otherwise is the one this paper argues they do not have.

KEY THESIS

The object of governance is not the model

It is the human-machine feedback loop: the evolving relationship among human judgment, machine assistance, environmental pressure, institutional purpose, and action over time.

This distinction is not academic, and it decides whether your programme can see anything that matters. Cognitive harm rarely originates in a single interaction. A prompt may be acceptable, an output accurate, an action authorised — while the *sequence* gradually alters how the human reasons, verifies, trusts, delegates and decides.

Governance fails when it is treated as telemetry. It succeeds when it is treated as discovery.

Temporal-linking is therefore a requirement, not an enhancement. Only by linking interactions across time can governance observe dependency drift, changes in verification behaviour, shifts in decision authority, and the quiet loss or strengthening of human agency.

Without it, a governance function is not merely incomplete. It is **structurally unable to detect its own primary failure mode**, and will continue reporting health for as long as the gates keep operating.

WHY EVENT-CENTRIC GOVERNANCE FAILS

Every point on this line passed

Monitoring isolated interactions confirms that each event was compliant, safe or authorised. It says nothing about whether the practice remains healthy — and the two diverge quietly, over quarters, in one direction.



- **Dependency drift develops gradually.** There is no day on which it happens, which means there is no day on which anyone could have escalated it.
- **Skills erode once verification becomes routine, then skipped.** The erosion is invisible while the outputs remain correct — and outputs remain correct right up until the case the system handles badly.
- **Automation bias rises as confidence is misplaced.** Your most experienced people are the most exposed, because their confidence is the least likely to be questioned.
- **Conditions shift underneath a stable control.** Yesterday's safe interaction becomes tomorrow's risk without anything in the record changing.

Event logs answer the wrong question. They report what happened at a point in time, not what the pattern of interactions is doing to the people who rely on it. **The sequence, not the single event, is the relevant unit of analysis** — and almost nothing in a modern governance stack is built to hold a sequence.

TEMPORAL-LINKING

Not logging. Association.

Temporal-linking connects interactions across time. It is not telemetry, and a data lake is not a substitute for it. It is the *contextual linking* of interactions so that governance can form associations between conditions and outcomes.

The system must learn, adapt and detect patterns across a time horizon on its own — because the humans inside the loop are, by construction, the least able to notice what is happening to them.

This is what makes governance diagnostic rather than retrospective. It produces leading indicators of risk and the ability to intervene before harm emerges, transforming governance from a ledger of transactions into an instrument of stewardship.

The unit of AI governance cannot remain the isolated interaction.

REDUCED PRACTICE

Operational control masquerading as stewardship

This is not a distinction between bad controls and good ideals. Controls remain necessary. Gates, approvals, thresholds, audit logs, risk ratings and enforcement points are all part of responsible governance. The failure begins when these instruments are mistaken for governance itself.

Reduced practice treats governance as the management of discrete objects and events. Is the system approved? Is the action authorised? Does a control exist? Did the output pass a threshold? Can a record be produced after the fact? Every question is legitimate. Every question is answerable at a point in time. That is precisely the limitation.

? Do users remain active observers, or have they become passive confirmers?

? Is confidence better calibrated, or merely inflated?

? Is verification improving, or disappearing?

? Is delegation preserving agency, or manufacturing dependency?

? Is the sequence of allowed events strengthening or weakening the capacity to steer?

Your programme cannot answer any of these. Not because it was built carelessly, but because it was built to inspect artifacts, and every one of these is a property of a trajectory.

AI does not merely produce outputs. It participates in the formation of judgment — framing options, reducing effort, accelerating decisions, shaping confidence, altering verification behaviour, and gradually changing what humans notice, trust, ignore and delegate. None of that is visible in a transaction. All of it is visible in a trajectory, to anyone holding an instrument that can see one.

Without that broader discipline, governance becomes **operational control masquerading as stewardship** — and the masquerade is convincing from the inside, because every indicator it produces is genuinely true.

GOVERNANCE AS PRAXIS

A living practice, not a compliance system

Praxis integrates theory, action, observation, reflection and correction. AI governance becomes praxis only when it observes how the human-machine relationship evolves over time, and adapts accordingly.

REDUCED PRACTICE ASKS

Was it allowed?

- Is the system approved?
- Was the action authorised?
- Does a control exist?
- Did the output pass?
- Can a record be produced?

PRAXIS ASKS

What is it doing to us?

- How are decisions being shaped over time?
- What dependencies are becoming invisible?
- How is the atmosphere of judgment changing?
- Is human agency being preserved or spent?
- Are we still able to steer?

Praxis learns from experience, adjusts to context, and strengthens human judgment through positive and negative feedback reinforcement — while steering development and use of AI in service of a shared mission and purpose.

THE REQUIREMENT

Governance must cultivate second-order observation: observ-

ing how observing is itself conditioned.

In practice this means steering through weather, ocean, vessel, human and desire simultaneously. It is not the operation of gates and flows. It is the art of navigating uncertainty with foresight and humility — and it requires instruments that no reduced practice has ever needed to build.

IMPLICATIONS

Govern delegation, not outputs

Control of individual gates, vessels and transactions is necessary but insufficient, because cognitive offloading and temporal effects diffuse decisive influence across time and context. Influence stops living where your controls are.

Effective governance must therefore preserve human agency, detect hidden dependency, and sustain the conditions for wise judgment. This is a shift from controlling flows to stewarding trajectories — and it is not a refinement of the current practice but a different instrument entirely.

Without temporal-linking, AI governance cannot observe its own effects. Without observing its own effects, it cannot be praxis.



CONCLUSION

Governance has always been both hydraulic and atmospheric. Reduced practice secured the gates and ignored the weather, the ocean, the human, and desire.

In the age of AI this myopia is dangerous, because influence now diffuses across time and context — into the very faculty that governance depends on to notice anything at all. The remedy is not more isolated controls. It is the restoration of governance as praxis: broader in scope, deeper in observation, oriented to trajectories.

AI changes the environment in which human judgment occurs. It does not merely answer questions or automate tasks. It shapes confidence, frames possibilities, reduces friction, alters attention, and changes what humans come to trust, question or ignore. That process is running in your organisation now, and has been for some time.

The discipline of governance must remain broader than the reduced practice that purports to represent it.

The relevant question was never whether the event was allowed.

It is what a sequence of allowed interactions does to judgment, agency, verification and responsibility over time.

SELECTED REFERENCES

- Wiener, N. (1948). *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT Press.
- von Foerster, H. (1992). Ethics and second-order cybernetics. *Cybernetics & Human Knowing*, 1(1), 9–19.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>
- Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
- Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776–778. <https://doi.org/10.1126/science.1207745>
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), Article 188. <https://doi.org/10.1145/3449287>

SOURCE NOTES AND CONCEPTS

This paper uses cybernetics as a steering metaphor and applies it to AI-enabled cognitive offloading. Wiener’s framing of control and communication anchors the first-order feedback concept. von Foerster’s second-order cybernetics motivates the claim that governance must observe the observing system, not merely the observed artifact. Clark and Chalmers, Sparrow et al., and Risko and Gilbert support the claim that cognition can be externally supported or offloaded; Parasuraman and Riley and Buçinca et al. support the concern that automation and AI-assisted decision-making can produce misuse, overreliance, and the need for cognitive forcing functions.

The block quotes are original working formulations developed for this document: “AI exposed the poverty of a winged (reactionary) governance stance”; “Cognitive harm is not in the event. It is in the trajectory”; “Operational control pretending to be governance is the core failure of reduced practice”; and “Without temporal-linking, AI governance cannot observe its own effects. Without observing its own effects, it cannot be praxis.”

The drawing is an original conceptual work. It depicts an unrealistically controlled simulation of reality as the current state of governance, contrasted with the reality of governance as a navigator steering through weather, ocean, vessel, human desire and trajectory. The dependency-drift chart is illustrative rather than empirical; it renders the paper’s central claim about sequence, not a measured dataset.