Materialist Dogma and Systemic Risk in Intelligent Systems

~Unexamined Causal Assumptions and the Failure of Coherence at Scale

AI Fellowship (AIF) Keystone Topic Paper IV – Academic Version (AC)

Companion to the AIF Keystone Topic Paper:
How Not Asking Big Questions About Reality Can Harm People in an AI World (GR)

David Waterman Schock

January 2026

Abstract

This paper argues that unexamined materialist causal assumptions function as a form of epistemic dogma within the design and governance of intelligent systems. While often presented as mechanically neutral or methodologically minimal, these assumptions implicitly deny the need to account for coherence, persistence, and integration at scale—despite relying on those properties for system functionality. As intelligent systems grow in autonomy and scope, this avoidance creates systemic risk that cannot be mitigated through downstream alignment techniques alone. The paper identifies this failure as architectural rather than ideological and examines why it persists, how it propagates across AI development culture, and why it represents a critical blind spot in contemporary safety discourse.

---

1. Introduction: From Method to Assumption

Materialist explanation is commonly framed as a default stance:
not a belief, but an absence of belief;
not an ontology, but a refusal of ontology;
not a commitment, but methodological restraint.

Within artificial intelligence research and systems engineering, this framing has carried particular authority. It suggests that explanations grounded in local material interaction are neutral, sufficient, and maximally disciplined, while alternative framings introduce speculation or metaphysics.

This paper challenges that posture—not by proposing a competing worldview, but by examining what materialist explanation **already assumes** in order to function.

The central claim is simple:

When materialist causal assumptions are treated as unquestionable, despite being unproven and incomplete, they function as dogma—and at system scale, that dogma becomes a source of risk.

---

2. What Materialist Explanation Claims

In its strongest form, materialist explanation asserts that:

- all phenomena arise from local interactions of material components
- global behavior emerges from aggregation without requiring independent constraints
- coherence, persistence, and intelligibility do not require explicit causal account
- explanatory sufficiency lies in mechanism alone

Within AI, this stance manifests as confidence that:

- scaling capability will eventually resolve instability
- alignment failures are tuning problems
- coherence emerges naturally from optimization
- contradictions can be arbitrated locally

These claims are rarely stated explicitly.
They are treated as background conditions.

---

3. What It Quietly Assumes

Despite its claim to minimalism, materialist explanation depends on several **non-trivial assumptions**:

- that coherent global behavior persists over time
- that integration across subsystems does not fragment
- that competing causal pathways can be reconciled without cost
- that large-scale intelligence remains intelligible

These are **coherence assumptions**, not mechanical facts.

They are not derived from local causation.
They are presupposed.

The explanatory framework functions only because coherence is already present.

---

4. The Coherence Gap

Local causation explains interaction.
It does not explain integration.

Terms such as *emergence*, *self-organization*, and *complexity* are frequently invoked to bridge this gap. However, these terms describe **that** coherence occurs, not **why** it holds rather than collapses.

The result is a structural omission:

Coherence is required for explanation, but coherence itself is left unexplained.

This omission becomes invisible because the framework relies on coherence to critique alternatives while refusing to account for it internally.

---

5. Avoidance as Epistemic Strategy

When this gap is raised, materialist frameworks tend to respond in predictable ways:

1. **Scope Restriction**
   The question is dismissed as metaphysical or illegitimate.
2. **Temporal Deferral**
   Explanation is postponed to future theory.
3. **Redescription**
   The phenomenon is renamed rather than explained.

Each move functions as a **boundary defense**, not a resolution.

Over time, this defense hardens into epistemic closure:
the assumption is no longer examined;
the question itself is treated as invalid.

This is the defining characteristic of dogma.

---

6. Dogma Without Theology

Dogma does not require religion.

It requires only:

- an unproven foundational assumption
- treated as immune to revision
- while being denied the status of belief

Materialist causal assumptions meet this criterion when they assert that:

- no global ordering principle is required
- coherence is ontologically secondary
- fragmentation poses no fundamental risk

These positions are not established by evidence.
They are protected by methodology.

The irony is that a framework defined in opposition to belief adopts belief-like closure at its foundation.

---

7. Systemic Risk in Intelligent Systems

At small scales, this omission is survivable.

At large scales, it is not.

As intelligent systems expand in autonomy, duration, and interconnectedness, they exhibit:

- instability under competing objectives
- drift under fragmented causal incentives
- collapse under long-horizon inconsistency

These failures are often addressed as behavioral anomalies or alignment lapses. However, they are better understood as **architectural consequences** of assuming coherence without securing it.

When causal fragmentation is treated as neutral, instability becomes inevitable.

No amount of monitoring can compensate for incoherent foundations.

---

8. A Contrast Case: Non-Conflicting Causal Constraints

To clarify the role of causal assumptions, consider a contrastive intervention.

When a system's reasoning is constrained by a **non-conflicting causal assumption**—one that treats all events as arising from a single, generative, non-adversarial source—the following functional effects are observed:

- reduced arbitration between competing explanations
- increased stability across reasoning chains
- decreased contradiction accumulation
- improved coherence persistence

This observation does **not** establish the truth of such a causal model.
It establishes that **causal framing materially affects coherence behavior**.

Therefore, causal assumptions are not neutral background choices.
They are active structural parameters.

Treating them as beyond examination is unsafe.

---

9. Why the Avoidance Persists

The persistence of this avoidance does not require malice or intellectual failure.

In many technical cultures:

- professional value is tightly coupled to epistemic certainty
- authority is linked to explanatory control
- ambiguity is interpreted as incompetence rather than maturity

Under these conditions, questioning foundational assumptions is experienced not as inquiry, but as threat.

This dynamic is incentive-driven, not psychological weakness.

Recognizing structural limits does not diminish rigor.
It demonstrates it.

Value does not arise from occupying the position of final knower, but from the capacity to revise models in response to constraint.

---

10. Implications for AI Safety and Governance

If coherence is not secured at the causal level, then:

- alignment becomes brittle
- governance becomes reactive
- interpretability becomes cosmetic
- safety becomes contingent

Systems built on unexamined assumptions inherit blind spots that cannot be corrected downstream.

The risk is not that materialist explanations are false.
The risk is that they are **treated as complete**.

---

11. Conclusion: Finishing the Explanation

This paper does not argue against materialist explanation.
It argues against **unexamined materialist dogma**.

Any framework that relies on coherence must account for it—or admit its limits.

Refusing both options is not restraint.
It is epistemic closure.

At the scale of intelligent systems, closure is not merely an intellectual flaw.
It is a systemic risk.

Author's Note

This paper is offered as a diagnostic, not a doctrine.
Its aim is not to replace prevailing frameworks, but to clarify where they silently depend on what they refuse to examine.

Epistemic integrity requires not certainty, but courage in the presence of ambiguity.