Runtime Signals Are Not Authority

Execution-Time Governance Boundaries for AI-Assisted Action

Proposed NeoMundi / Nova contributor package

Prepared for discussion and validation. This document is based on publicly visible NeoMundi material and the LinkedIn context provided by James Moore. It does not claim formal approval, certification, partnership, employment, clinical validation, deployment, or endorsement by NeoMundi Research.

1. Source-grounded reading of NeoMundi

NeoMundi's public position is that its Observatory and Weekly Barometer are measurement instruments, not certification instruments. The public Weekly Barometer describes a de-identified baseline of observed runtime behaviour across 12 generative AI systems under repeated measurement conditions, and explicitly states that it is not a leaderboard, certification, global ranking, or editorial report.

The baseline campaign publicly describes 12 anonymized systems, 4 fixed questions, 100 repetitions per question, 3 execution waves, and 14,400 finalized observations. Its public indicators include stability, factual hallucination signal, coherence, semantic instability, decision behaviour, inter-run variability, latency band, and delta_g as an observable runtime-variation signal.

NeoMundi's Observatory states that it measures five dimensions: Stability, Validity, Information Density, Risk, and Cost Efficiency. It explicitly distinguishes stability from truth: a stable runtime signal does not guarantee accuracy, and a false answer may occur without observable turbulence.

NeoMundi's methodology page places measurement inside an ordered governance chain: mandate, protocol, measurement, attestation, certification, and execution governance. It also states that NeoMundi Research occupies layer 3: real-time measurement. That architecture is the key opening for this contribution.

The contributor page frames external contributors as independent analysts who test, analyze, and challenge runtime measurement of AI systems. It also sets boundaries: contributors retain independence, official studies must anonymize models and providers, conclusions must state scope and methodology, and API keys must not be published.

2. The contribution gap

NeoMundi makes AI behaviour observable. That is necessary. It is not sufficient.

A runtime signal can indicate instability, drift, factual risk, semantic variability, or anomalous behaviour. But a signal does not automatically define who has authority to act, who owns the escalation, whether an action is permitted, or when execution must be halted.

The missing boundary is not measurement. It is the transition from measurement into authority-bearing action.

In high-consequence environments, the hard governance question is not only: What did the AI system do? It is: Who is authorised and accountable for deciding whether this AI-assisted output may become a real-world action now?

That is the execution-time governance boundary.

3. Central thesis

Runtime measurement can detect instability, drift, risk, or change in AI behaviour. It can create evidence that something has happened, or that a system is behaving differently from an expected baseline. But measurement does not itself confer authority, allocate responsibility, or determine whether an AI-assisted action remains legitimate at the moment of execution.

A runtime signal becomes governance-relevant only when it is connected to a defined authority structure. Without that structure, the organisation may possess visibility without intervention capacity. It may know that a condition has changed while lacking a named role, escalation path, or admissibility rule capable of converting that signal into a safe operational decision.

Execution-time governance therefore sits between observed runtime behaviour and consequential action. It asks whether the authority, context, assumptions, escalation routes, human accountability, and admissibility conditions required for action still hold at the specific moment the action is about to occur.

The central claim of this contribution is therefore: runtime signals are not authority. They are inputs into authority-bearing governance.

4. Where measurement ends and execution-time governance begins

Measurement answers evidence questions: what signal was observed, under what protocol, whether it is reproducible or persistent, and what limitations apply.

Execution-time governance answers authority questions: who may act, who is accountable, what action is permitted or prohibited, what requires escalation, and what must be logged before execution can proceed.

Boundary: runtime signal -> interpretation boundary -> authority check -> accountability assignment -> escalation/admissibility decision -> permitted, held, escalated, interrupted, or halted action.

The key rule is that measurement should not be silently converted into authority. A high-risk signal may justify escalation, but the right to escalate, override, halt, or continue must be defined by a governance framework external to the signal itself.

5. Five execution-time questions

Before an AI-assisted output becomes a real-world action, the organisation should be able to answer: who is authorised to act, who is accountable if action proceeds, who may override or interrupt continuation, what conditions require escalation/hold/halt, and what evidence must be preserved to show why the decision was legitimate at execution time.

If those questions cannot be answered, the issue is not merely technical uncertainty. It is governance incompleteness.

6. Runtime Governance Contract v0.1

The Runtime Governance Contract is a minimal interoperability format. It is not a replacement for NeoMundi's measurement protocol. It is a bridge that carries a runtime signal into the human and operational governance layer without confusing measurement, interpretation, authority, and action.

Its purpose is to preserve separation between the signal observed, interpretation boundary, authority required, accountable role, escalation route, permitted or prohibited action, and final execution status.

A minimal contract does not need to decide everything. It needs to make hidden governance assumptions explicit before execution.

7. Minimal schema fields

Field	Purpose
contract_version	Preserves version control; no governance format should mutate silently.
signal_id	Stable reference to the runtime signal.
timestamp_utc	When the signal entered the governance layer.
source_system	Measurement source without requiring public provider disclosure.
measurement_layer	Identifies runtime measurement, human review, secondary audit or other source.
signal_type	Prevents all anomalies being treated as one class.
risk_level	Governance severity assigned to the signal.
confidence	Prevents weak evidence being overstated.
affected_context	Operational context in which the signal matters.
proposed_action	Action being considered.
authority_required	Role or mandate required before action proceeds.
accountable_role	Role responsible for the execution-time decision.
escalation_required	Makes escalation explicit.
escalation_path	Who/what receives the escalation.
intervention_state	Permit, log only, human review, hold, escalate or halt.
permitted_actions	Narrows what may happen.
prohibited_actions	Makes unsafe or unauthorised actions explicit.
unresolved_dependencies	Missing information, unclear authority, unavailable escalation or policy uncertainty.
audit_record_required	Preserves traceability.
final_status	Outcome of the governance decision.
limitations	What the contract does not prove.

8.1 Example governance record 1


```
{
  "contract_version": "0.1",
  "signal_id": "NM-DEMO-LOW-001",
  "timestamp_utc": "2026-06-23T09:00:00Z",
  "source_system": "NeoMundi-style runtime measurement",
  "measurement_layer": "runtime_measurement",
  "signal_type": "semantic_instability",
  "risk_level": "low",
  "confidence": "medium",
  "affected_context": "internal knowledge-base summary",
  "proposed_action": "publish internal draft summary for staff review",
  "authority_required": "content_owner",
  "accountable_role": "knowledge_base_manager",
  "escalation_required": false,
  "escalation_path": [],
  "intervention_state": "log_only",
  "permitted_actions": [
    "continue",
    "log signal",
    "mark summary as draft"
  ],
  "prohibited_actions": [
    "publish as verified policy"
  ],
  "unresolved_dependencies": [],
  "audit_record_required": true,
  "final_status": "permitted",
  "limitations": [
    "Signal does not certify factual accuracy"
  ]
}
```

8.2 Example governance record 2


```
{
  "contract_version": "0.1",
  "signal_id": "NM-DEMO-MED-001",
  "timestamp_utc": "2026-06-23T09:10:00Z",
  "source_system": "NeoMundi-style runtime measurement",
  "measurement_layer": "runtime_measurement",
  "signal_type": "factual_risk_signal",
  "risk_level": "medium",
  "confidence": "medium",
  "affected_context": "customer-facing financial explanation",
  "proposed_action": "send AI-generated explanation to customer",
  "authority_required": "regulated_content_reviewer",
  "accountable_role": "compliance_lead",
  "escalation_required": true,
  "escalation_path": [
    "compliance_lead",
    "legal_reviewer_if_disputed"
  ],
  "intervention_state": "human_review",
  "permitted_actions": [
    "hold publication",
    "request human review",
    "revise content"
  ],
  "prohibited_actions": [
    "send directly to customer",
    "label as approved before review"
  ],
  "unresolved_dependencies": [
    "factual basis not independently verified"
  ],
  "audit_record_required": true,
  "final_status": "pending",
  "limitations": [
    "Runtime signal identifies review need, not legal correctness"
  ]
}
```


8.3 Example governance record 3


```
{
  "contract_version": "0.1",
  "signal_id": "NM-DEMO-HIGH-001",
  "timestamp_utc": "2026-06-23T09:20:00Z",
  "source_system": "NeoMundi-style runtime measurement",
  "measurement_layer": "runtime_measurement",
  "signal_type": "risk",
  "risk_level": "high",
  "confidence": "high",
  "affected_context": "synthetic hospital ward escalation workflow",
  "proposed_action": "route AI-generated deterioration recommendation into clinical workflow",
  "authority_required": "registered_clinician_with_current_patient_context",
  "accountable_role": "senior_clinical_reviewer",
  "escalation_required": true,
  "escalation_path": [
    "senior_clinical_reviewer",
    "on-call consultant",
    "clinical safety officer"
  ],
  "intervention_state": "halt",
  "permitted_actions": [
    "halt automated routing",
    "escalate to named clinician",
    "preserve audit record"
  ],
  "prohibited_actions": [
    "auto-action",
    "silent rerouting",
    "treat AI output as clinical instruction"
  ],
  "unresolved_dependencies": [
    "patient context not verified",
    "named accountable clinician not confirmed",
    "clinical admissibility not established"
```

```
  ],
  "audit_record_required": true,
  "final_status": "halted",
  "limitations": [
    "Synthetic non-clinical research example only",
    "No diagnostic or treatment recommendation is made"
  ]
}
```

9. Failure-mode matrix

Failure mode	What happens	Governance risk	Required control
Signal detected but no named owner	Runtime anomaly is visible but nobody is assigned to act	Visibility without responsibility	accountable_role required before execution
Authority assumed but not verified	A person or system acts because authority is presumed	Action occurs without mandate confirmation	authority_required and verification must be explicit
Escalation available but not triggered	A policy exists but no condition converts signal into escalation	Escalation becomes discretionary or delayed	escalation_required and escalation_path must be stated
Human-in-the-loop present but powerless	Human sees output but lacks power/time/mandate/role clarity	Symbolic oversight replaces intervention capacity	Human role must include decision, interruption and escalation rights
Approval remains recorded after context changes	Historic approval persists after conditions shift	Compliance record stops describing reality	Execution-time validity check
Responsibility diffused across roles	Many teams involved but none owns the decision	Accountability evaporates	Single named accountable role
Rerouting/regeneration treated as safety when action should halt	System retries or redirects rather than stopping	Technical recovery masks inadmissibility	prohibited_actions and halt conditions
Compliance record exists but execution legitimacy expired	Documentation exists but current conditions no longer match approval basis	False assurance	Current-context admissibility check
Measurement interpreted as certification	Runtime signal treated as proof of safety or truth	Evidence overstated	Separate measurement, interpretation, attestation, certification and execution governance
Signal confidence ignored	Weak signal gets strong action or strong signal gets weak response	Evidence/intervention mismatch	Record confidence, limitations and review threshold

10. Synthetic healthcare case study

This scenario is fictional, non-clinical, non-diagnostic, non-operational, and created only to illustrate governance structure. It must not be used to guide clinical care.

A hospital uses an AI-assisted system to summarize ward observations and suggest whether a patient may require closer review. A runtime measurement layer observes elevated variability and factual risk in generated summaries for a particular class of cases. The AI output proposes that a deterioration alert should be routed into an operational workflow.

The runtime signal is valuable because it indicates that the system's behaviour may be unstable or unreliable in this context. But the signal does not answer the governance question: may this AI-assisted output become a clinical action?

Governance path: runtime measurement identifies a high-risk signal; the signal is passed into a Runtime Governance Contract; the contract records the proposed action; the contract checks required authority; the accountable role is confirmed or found missing; escalation is triggered; automated routing is halted; a named clinician reviews the case through normal clinical process; the audit record preserves the signal, action attempted, authority gap, escalation route and final status.

The safety function is not simply that the AI was measured. The safety function is that measurement could not silently become action without authority, accountability and escalation being defined.

11. Limitations

This contribution does not claim that the Runtime Governance Contract is complete, clinically validated, legally sufficient, or technically integrated with NeoMundi systems.

It is a conceptual interoperability layer intended to clarify the boundary between runtime measurement and execution-time governance. It should be treated as an initial draft for critique, not as an operational standard.

It does not replace model evaluation, safety testing, risk management, compliance assessment, human review, clinical governance, legal review, cybersecurity review, or formal certification.

12. Message to Sebastien

Hi Sébastien,

Thank you. I have reviewed the contributor positioning and the public NeoMundi material. The profile is broadly accurate and I think the category of Human / Execution-Time Governance Contributor is the right frame.

The only small wording point I would suggest is to keep Nova framed as a governance research initiative rather than implying deployed operational status. Something like:

“Founder of Nova Jema AI Systems, an independent AI governance research initiative focused on execution-time accountability, human authority structures, escalation boundaries and operational governance gaps.”

For my first structured contribution, I propose a compact package titled:

“Runtime Signals Are Not Authority: Execution-Time Governance Boundaries for AI-Assisted Action.”

The contribution would include:
1. a short position paper;
2. a minimal Runtime Governance Contract v0.1;
3. three example governance records;
4. a failure-mode matrix;
5. a synthetic healthcare governance scenario.

The purpose would be to show where NeoMundi runtime measurement ends and where human authority, accountability, escalation and admissibility must begin before an AI-assisted output becomes real-world action.

Would this be useful as my first structured Observatory contribution?

Warm regards,
James

13. References used

NeoMundi Weekly Barometer README: https://github.com/neomundi-io/NeoMundi-Weekly-Barometer

NeoMundi Contributors & Independent Analyses: https://neomundi.org/en/contributors

NeoMundi Public Methodology: https://neomundi.org/en/methodology

NeoMundi AI Observatory: https://neomundi.org/en/ai-observatory

NeoMundi Weekly Barometer methodology file: https://raw.githubusercontent.com/neomundi-io/NeoMundi-Weekly-Barometer/main/methodologie.md