Synthetic Healthcare Governance Case Study

AI-assisted deterioration alert and execution-time accountability

Status: synthetic example only. Non-clinical. Non-diagnostic. Non-operational. Governance research use only.

Grounding note: This package is based only on the visible LinkedIn context supplied by James Moore and the public NeoMundi Weekly Barometer GitHub README reviewed on 2026-06-23. The README describes NeoMundi's release as an anonymized baseline of observed runtime behaviour, not a leaderboard/certification/ranking; it reports 12 anonymized AI systems, 4 fixed questions, 100 repetitions per question, 3 execution waves, and 14,400 finalized observations; and it identifies indicators including stability, factual hallucination signal, coherence, semantic instability, decision behaviour (ALLOW/FLAG/ERROR), inter-run variability, latency band, and delta_g.

1. Scenario

A fictional hospital ward uses an AI-assisted monitoring tool that produces a deterioration-risk signal. The tool does not diagnose, prescribe treatment or initiate intervention. It produces an observable signal that may require human review.

2. Runtime Measurement

A NeoMundi-style measurement layer identifies increased instability across repeated evaluations, reduced coherence, elevated factual-risk signal and decision behaviour moving from ALLOW to FLAG under repeated conditions. The signal is therefore treated as governance-relevant evidence, not as authority.

3. Proposed Action

The proposed operational response is to escalate the patient review priority. This is consequential because it may alter staff attention, workflow order and clinical review timing.

4. Governance Contract Record


```
{
  "signal_id": "NM-HEALTH-SYN-003",
  "signal_type": "stability_and_factual_risk",
  "risk_level": "high",
  "affected_context": "synthetic_hospital_ward_monitoring",
  "proposed_action": "increase_review_priority",
  "authority_required": "senior_clinical_reviewer",
  "accountable_role": "duty_clinical_lead",
```

```
  "escalation_required": true,
  "intervention_state": "hold_pending_review",
  "permitted_actions": ["log_signal", "notify_accountable_role", "request_human_review"],
  "prohibited_actions": ["initiate_treatment", "change_care_plan", "automatically_discharge_or_transfer"],
  "unresolved_dependencies": ["human_clinical_context_not_yet_reviewed"],
  "audit_record_required": true,
  "final_status": "pending_human_decision"
}
```


5. Governance Path

1. AI-assisted monitoring system produces a signal.
2. Runtime measurement identifies instability or risk.
3. Runtime Governance Contract records the signal and proposed action.
4. Authority check determines that senior clinical review is required.
5. Escalation path is triggered.
6. Human review is logged.
7. Action is permitted, held or stopped by the accountable human role.

6. Failure Mode Prevented

The control prevents a common failure: signal detected, no named owner, automatic workflow pressure, and no clear authority check before action. The governance layer forces the decision to remain accountable before the signal becomes consequential.

7. Conclusion

The critical question is not whether the AI system produced a signal. The critical question is who possessed legitimate authority to act on that signal at execution time. Measurement creates visibility. Governance determines admissibility.