

Operational Pathology of AI-Augmented Work

This document maps public AI and automation failures into the Epistema / AIWI lifecycle vocabulary. It is not a forensic or legal analysis of the incidents. It is an interpretive architecture note: the purpose is to test whether ACTECRU and the Workflow Value Ledger explain recognizable failure patterns in AI-augmented work.

AI fails in organizations when outputs become commitments before the workflow is ready.

Signal was produced -> Evidence was weak or missing -> Commitment was premature or unbounded -> Outcome was negative -> Calibration lagged or never occurred

Compact Reference Mapping

Case	Lifecycle Stage	Missing Transition	Observable Symptom	WVL Evidence
Air Canada	Evaluation -> Commitment	Policy validation before customer commitment	Unable to reconstruct authoritative source	Missing evidence provenance; missing commitment authorization
Samsung	Admission -> Classification	Sensitive information classification	Irreversible external disclosure	Missing admission record; missing data classification
Zillow	Outcome -> Update	Drift-driven recalibration	Continued execution despite mounting losses	Drift signal ignored; calibration absent
Bing (Sydney)	Transformation <-> Evaluation	Continuous evaluation during long reasoning trajectories	Behavioral instability	No intermediate evaluation checkpoints
Tokenmaxxing	Evaluation -> Outcome	Cost-to-value assessment	Rising spend with flat organizational value	Cost provenance disconnected from workflow quality

Case Interpretations

1. Air Canada: Commitment Without Evaluation

The chatbot answer crossed from informational output into customer-facing commitment. The Workflow Value Ledger should show source policy, validation status, customer-facing commitment status, and accountable actor or system boundary.

2. Samsung: Admission and Classification Failure

The workflow needed to know what was allowed to enter. The ledger should show admission status, data classification before model/tool use, external-service boundary, and audit trail.

3. Zillow Offers: Automation Governance Under Drift

The governance issue is continuing to execute when the environment has changed faster than the governance loop. The ledger should show outcome signals, drift evidence, calibration records, and changed admission thresholds.

4. Bing / Sydney: Session-Level Trajectory Governance

The governance object is not only an output. It is the trajectory. The ledger should show turn count, drift indicators, intermediate evaluation checkpoints, constraint refresh events, and termination thresholds.

5. Tokenmaxxing: Cost Without Value Alignment

Cost accounting alone does not answer the governance question. The relevant question is: what did the spend produce, and was the result worth committing to?

Why This Matters for Epistema

The cases are deliberately different. They show that AI governance cannot be reduced to model evaluation, policy documentation, risk registers, cost dashboards, or after-the-fact incident review. Epistema's narrower structural claim is that AI-augmented work needs observable lifecycle state.

The Workflow Value Ledger is the compression layer for that state. It joins cost provenance, admitted inputs, evidence coverage, risk shape, commitment status, human gate events, resolution status, and calibration records.

The ledger is not cost accounting and not traceability alone. It is the structural coupling of resource evidence with decision-quality governance.

Public-Language Boundary

Prefer: AI fails in organizations when outputs become commitments before the workflow is ready. Avoid overclaiming formulations such as: Every major AI failure is a lifecycle mismatch.

Sources

* The Guardian: Air Canada chatbot liability

* Bloomberg and TechCrunch: Samsung generative-AI restriction

* Axios: Zillow Offers shutdown

* CNN, Bloomberg, and The New York Times: Bing / Sydney long-session behavior

* Business Insider and arXiv: tokenmaxxing and AI cost discipline