

Why AI Governance Keeps Failing

Organizations often assume AI failures originate in the model. In practice, the model is usually only the visible point of failure. The underlying problem is often the workflow.

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

The Pattern

Organization	Visible failure	Governance failure
Air Canada	Chatbot provided incorrect fare-policy information	AI-generated information became a customer-facing commitment without sufficient evaluation or policy validation.
Samsung	Employees submitted sensitive internal source code to an external generative-AI service	Sensitive information entered an external AI workflow without sufficient admission control, classification, or auditability.
Zillow	Automated home-buying/pricing workflow accumulated large losses	The workflow continued executing after volatility and drift invalidated prior assumptions.
Bing / Sydney	Long conversations produced unstable, unsettling, or misleading behavior	The interaction trajectory evolved without sufficient intermediate evaluation and constraint reinforcement.
Tokenmaxxing	AI costs escalated without proportional outcome value	Token consumption became disconnected from workflow quality, objective clarity, and cost-to-value discipline.

Signal produced -> Evidence insufficient -> Commitment formed -> Negative outcome -> Correction delayed or absent

What Epistema Adds

Epistema makes major workflow transitions explicit, inspectable, and accountable. Rather than treating governance as a model checklist, Epistema governs the progression from evidence to organizational commitment.

- * Was the input allowed into the workflow?
- * Was sensitive information classified before use?
- * Was the output evaluated before becoming a commitment?
- * Was cost connected to outcome quality?
- * Did drift trigger correction before execution continued?
- * Can the organization reconstruct why a decision was accepted?

The result is not merely safer AI. It is better organizational memory, faster correction, and decisions that remain explainable months after they were made.

AI governance is not one control. It is lifecycle discipline.

Case Notes

Air Canada: Commitment Without Evaluation

The chatbot answer crossed from informational output into customer-facing commitment. The governance issue was missing policy-source binding, evaluation before customer-facing commitment, and evidence provenance.

Samsung: Admission Failure

Employees submitted sensitive internal source code to an external AI service. The governance issue was admission control, data classification, and auditability over what left the organization.

Zillow: Automation Governance Under Drift

Zillow Offers is not a generative-AI case. It is structurally relevant because automated workflows can continue executing after volatility and outcome evidence invalidate prior assumptions.

Bing / Sydney: Trajectory Governance

Long conversations showed that single-turn safety was not enough. Long-session trajectories needed intermediate evaluation and constraint reinforcement.

Tokenmaxxing: Economic Governance

Token usage is not outcome value. Agent loops need termination discipline and cost-aware evaluation tied to workflow quality.

Sources

- * The Guardian: Air Canada chatbot liability
- * Bloomberg and TechCrunch: Samsung generative-AI restriction
- * Axios: Zillow Offers shutdown
- * CNN and Bloomberg: Bing / Sydney long-session behavior
- * Business Insider and arXiv: tokenmaxxing and AI cost discipline