AI regulation without theater

AI Regulation Without Theater

A policy document can describe responsible artificial intelligence in immaculate language and still fail to change a single unsafe operation. It may require transparency without defining what an operator must preserve, promise human control without granting humans the authority to stop deployment, or demand an impact assessment no one can reproduce. Regulation then becomes a performance rather than a control system.

Effective governance begins at the point where a decision can be observed and changed. Every consequential AI workflow has an owner, an explicit purpose, authorised inputs and outputs, a record of model and prompt versions, tests for known failures and an escalation route. The organization decides which actions remain advisory and which require human approval or hard technical restrictions.

The evaluation must match the service. A language assistant preparing a public notice has different failure modes from a system ranking clinical appointments, controlling an industrial process or recommending a benefit decision. Technical validation therefore combines domain-specific scenarios, adversarial testing, privacy and discrimination review, operational monitoring and post-deployment incident handling.

Procurement has a central role. Contracts should preserve access to logs, model-change notifications, exportable records, security boundaries, responsibility for incidents and termination rights. A supplier promise is useful, but an acceptance test that the customer can rerun is better. The strongest regulation makes misleading performance and hidden dependencies difficult to sustain.

The public should judge AI governance by the behaviour of deployed systems and the ability to correct them. If a rule cannot be translated into an inspectable responsibility, a test, a technical control or an enforceable remedy, it is not yet an adequate operational safeguard.